Factoring with Cyclotomic Polynomials*

By Eric Bach and Jeffrey Shallit

Dedicated to Daniel Shanks

Abstract. This paper discusses some new integer factoring methods involving cyclotomic polynomials.

There are several polynomials $f(X)$ known to have the following property: given a multiple of $f(p)$, we can quickly split any composite number that has p as a prime divisor. For example—taking $f(X)$ to be $X - 1$—a multiple of $p - 1$ will suffice to easily factor any multiple of p, using an algorithm of Pollard. Other methods (due to Guy, Williams, and Judd) make use of $X + 1$, $X^2 + 1$, and $X^2 \pm X + 1$.

We show that one may take f to be Φ_k, the kth cyclotomic polynomial. In contrast to the ad hoc methods used previously, we give a universal construction based on algebraic number theory that subsumes all the above results. Assuming generalized Riemann hypotheses, the expected time to factor N (given a multiple E of $\Phi_k(p)$) is bounded by a polynomial in k, $\log E$, and $\log N$.

1. Introduction. This paper discusses a new method for factorization of numbers, given partial information about the factors. Our algorithm includes some well-known methods of factoring as special cases, and provides a synoptic view of a large class of factorization methods.

Let N be a composite number with unknown factorization, and let p be an unknown prime dividing N.

Several investigators have observed that it is easy to split N if one knows *any* multiple of $p - 1$. Miller used this observation to show that computing the Euler ϕ-function is equivalent in difficulty to factoring N [17]. Based on this idea, Pollard devised a "$p - 1$" method of factoring [19]; it was also known to D. N. and D. H. Lehmer [23]. This method is "often spectacularly successful since it can sometimes find a quite enormous factor p with very little computing if $p - 1$ splits entirely or almost entirely into a product of small primes." [4]

Similarly, it is easy to split N if one knows a multiple of $p + 1$; this was pointed out by Guy in [8] and made the basis of recent implementations of Williams [23] and Brent [4]. Williams and Judd ([24], [25]) extended these ideas to the polynomials $p^2 + 1$ and $p^2 \pm p + 1$; their methods, although originally designed for prime testing, easily extend to factoring.

It is natural, given these facts, to make the following conjecture: since $X \pm 1$, $X^2 + 1$, and $X^2 \pm X + 1$ are all defining polynomials for complex roots of unity, the above results should be provable using special cases of some unknown "cyclotomic" factoring method.

Further support for this conjecture is provided by [1], in which we (along with Miller) showed that finding $\sigma(N)$, the sum of N's divisors, is equivalent in difficulty to factoring N. Since σ is expressible as a product of cyclotomic polynomials, our results also showed that one can quickly split N using a multiple of $\Phi_k(p)$, where $\Phi_k(X)$ is the kth cyclotomic polynomial and k is a prime power. Extending some ideas of Williams [25], we recently showed this to be true also if k has at most two prime factors [2].

In this paper we give a universal construction that implies *all* the above results and (assuming the Extended Riemann Hypothesis) proves the conjecture in the following substantive form: given any multiple of $\Phi_k(p)$, one can remove a factor p from N in random polynomial time.

We call an integer *B-smooth* if all of its prime factors are less than or equal to B. The methods of this paper give an algorithm with running time polynomial in B, k, and $\log N$ to extract a prime divisor p of N provided that $\Phi_k(p)$ is B-smooth. It should be noted that such primes p are probably rare, and so our results are unlikely to have great impact on the *practice* of factoring. However, we find this theoretically interesting for the reasons given below.

First, the previous methods involved many ad hoc techniques, using linear recurrences, finite fields, and so on. From a theoretical standpoint, it is interesting that all of these results can be explained as special cases of one algorithm.

Second, our results shed light on the paradigms that are used in designing factoring algorithms. Many such algorithms involve "pushing elements into subgroups" (see, e.g., [5]); our results give explicitly a large class of groups that lead to factoring algorithms.**

Finally, the use of generalized reciprocity laws in algorithms is just beginning to be explored. The analysis of our method makes essential use of these; in particular, we use a simple "higher reciprocity law" that can be stated with a minimum of extra concepts. Surely, this will find other uses in the construction of algorithms.

The rest of this paper is organized as follows. In Section 2, we discuss a "$p + 1$" method; our viewpoint caries over to arbitrary "cyclotomic" methods, presented in Sections 3 and 4. The running time is discussed in Sections 5–7. Finally, in Section 8 we present theorems on factorization.

We chose this division of topics for the following reason. Knowing basic field theory, one can read the first three sections and implement what might be called the "typical" version of the algorithm. To understand all the details of the algorithm, as well as a heuristic argument for the running time, one needs algebraic number theory at the level of [16]. Finally, to get the polynomial time bound, one has to know the basic results of class field theory, as presented in the appendix of [27].

** Just what groups can be used is an open question. There are factoring algorithms involving quadratic-field class groups [20] and elliptic curves [15], and it would be of interest to find some point of view that includes these, as well.

2. The $p+1$ Method. Before presenting our general construction, we discuss the $p+1$ factoring method. This was originally defined with linear recurrences; we will use algebraic number theory, as this gives an algorithm that is easier to generalize. Since we will give detailed proofs later, we will restrict ourselves to the case where N is a product of two odd primes p and q.

THE $p+1$ FACTORING ALGORITHM.
> Input: $N = pq$, and E, a multiple of $p+1$.
> Pick integers a, b, and d at random, and let $t = a + b\sqrt{d}$.
> By rationalizing the denominator, evaluate $x = \bar{t}/t \bmod N$.
> Compute $x^E \equiv u + v\sqrt{d} \bmod N$.
> Hope that $\gcd(u-1, v, n)$ splits N.

Why is this likely to work? Consider what is happening algebraically. Assuming that the integer d is not a perfect square (this is likely), t belongs to the quadratic extension field $\mathbf{Q}(\sqrt{d})$. By its construction, the norm of x—that is, the product of itself with its conjugate—will be 1.

We do all the computations in the ring $\mathbf{Z}[\sqrt{d}] \bmod N$, and by the Chinese Remainder Theorem for commutative rings ([12, p. 63]) this is isomorphic to the direct sum

$$(1) \qquad\qquad \mathbf{Z}[\sqrt{d}] \bmod p \oplus \mathbf{Z}[\sqrt{d}] \bmod q.$$

We further hope that d is not a square modulo p, making the first factor a finite field $\mathbf{F}_{p^2}$.

It is shown in the theory of finite fields that the conjugation map on a finite field of order p^2 is the *Frobenius automorphism* $x \to x^p$. Therefore, by raising to a power that is some multiple of $p+1$, we will get a power of the norm of x (which has to be 1) in the first direct summand of (1).

The above discussion has not involved q, and it seems likely that in the other piece of (1), nothing particularly noteworthy will happen. Put another way, unless we are unlucky, we will have the relations***

$$u - 1 \equiv v \equiv 0 \pmod{p}, \qquad u - 1, v \not\equiv 0 \pmod{q},$$

allowing the last step to remove the factor p.

We now retrace the above discussion, using a point of view that we will adhere to for the remainder of the paper.

By choosing d, we have selected a field K of degree 2, along with a generator σ for its Galois group (which is cyclic of order 2). In this case, the action of σ is given by

$$(x + y\sqrt{d})^\sigma = x - y\sqrt{d},$$

***If the second condition fails, p and q will appear together in the gcd, but they can usually be separated by lowering the value of E. The details are similar to those of the $p-1$ method; see [1].

and it is clear how to interpret $(x + y\sqrt{d})^{\sigma^i}$. We extend this notation to include *symbolic powers* whose exponents are polynomials in σ, with integer coefficients:

$$\alpha^{a_n\sigma^n+\cdots+a_1\sigma+a_0} \quad \text{denotes} \quad (\alpha^{a_n})^{\sigma^n}\cdots(\alpha^{a_1})^\sigma(\alpha^{a_0}).$$

We hope that modulo p, σ is the Frobenius map on a finite field of p^2 elements, for then

$$x^{p+1} \equiv t^{(\sigma-1)(\sigma+1)} \equiv t^{\sigma^2-1} \equiv 1 \pmod{p},$$

allowing us to remove p from N.

3. The $\Phi_k(p)$ Method. Let Φ_k be the kth cyclotomic polynomial; this is the unique monic polynomial whose roots are the primitive kth roots of unity. It can be shown ([12, p. 206]) that Φ_k has degree $\phi(k)$ and integer coefficients. We let

$$\Psi(X) = (X^k - 1)/\Phi_k(X);$$

this will also have integral coefficients.

This section discusses how to quickly remove a factor p from N, when we know a multiple E of $\Phi_k(p)$.

All the computations will be done in a certain finite ring (denoted R_m) which depends on an auxiliary prime m. This ring has the following properties. First, R_m is a free $\mathbf{Z}_N$-module of rank k with basis $\{\beta_0, \ldots, \beta_{k-1}\}$; this means that every member of R_m is uniquely a linear combination of the β_i's with integer coefficients taken modulo N. As part of our implementation of R_m, we will provide "multiplication tables", which are (implicit or explicit) matrices for the linear maps $x \to \beta_i x$. Second, R_m has an automorphism σ of order k; again, this is just a linear transformation. Finally, the multiplicative identity of R_m has some representation $1 = c_0\beta_0 + \cdots + c_{k-1}\beta_{k-1}$.

We will give details in the next section; for now, the reader is urged to keep the following "typical" case in mind. R_m is the polynomial ring $\mathbf{Z}[X]/(f_m(X), N)$, where f_m generates a cyclic field K_m of degree k, σ is a generator for the Galois group of K_m, and the basis $\{\beta_i\}$ consists of powers of X.

> THE $\Phi_k(p)$ FACTORING ALGORITHM.
>> Input: N (to be factored) and E (a multiple of $\Phi_k(p)$).
>> Repeat until N splits:
>>> Choose a prime $m \equiv 1$ modulo k.
>>> Construct the ring R_m, with automorphism σ.
>>> Choose $t \in R_m$ at random, and set $x = t^E$.
>>> For each i, $1 \leq i \leq k$, that is relatively prime to k:
>>>> Set $\tau = \sigma^i$.
>>>> Compute $y = x^{\Psi(\tau)}$.
>>>> Set $D = \gcd(c_i - y_i)$.
>>>> Try to split N with $\gcd(D, N)$.

What we hope for (in the typical case above) is this: f_m will be irreducible modulo some prime p dividing N, so R_m will have a direct factor isomorphic to the finite field $\mathbf{F}_{p^k}$. Then σ will have some power τ that is the Frobenius map $t \to t^p$ on this finite field. Then modulo p,

$$x^{E\cdot\Psi(\tau)} \equiv x^{\Phi_k(p)\cdot\Psi(p)} \equiv x^{p^k-1} \equiv 1,$$

and we further hope that this will *not* be congruent to 1 modulo some other divisor of N, allowing the last step to split N. (If the latter condition fails and $x^{E \cdot \Psi(\tau)} \equiv 1$ (mod N), a smaller value of E is likely to split N.)

(The reader may wish to glance at the appendix, which contains three examples.)

It is not obvious that the algorithm will work in polynomial time, but we can sketch the ideas involved in the typical case without too many technicalities.

For a prime m congruent to 1 mod k, let K_m be the unique field of degree k contained in the cyclotomic field generated by the mth roots of unity (the existence of this field follows from Galois theory). Further, let K_m be generated by a root of the monic integral polynomial $f_m(X)$. Then, for most primes p, $f_m(X)$ splits modulo p in the same way as $X^k - p$ splits modulo m. (This is a simple "higher reciprocity law"; we prove it later as Lemma 10.) Using this last fact, it suffices to find a prime $m \equiv 1$ (mod k) modulo which $X^k - p$ is irreducible. Among primes congruent to 1 mod k, Chebotarev's density theorem ([13, p. 169]) implies that the ones we seek have density $1/\phi(k)$, so heuristically at least, we expect to have to try $\phi(k)$ values of m before splitting N.

However, this heuristic argument neglects two things. We are interested in finding a *small m* that works, and we also want some assurance that the resulting factorization is nontrivial. We attend to these matters by analyzing a version of the algorithm that uses m in order from the sequence $k+1, 2k+1, 3k+1, \ldots$. We show below—assuming ERH—that the least m that is likely to lead to a nontrivial factorization is bounded by a polynomial in k and $\log N$. This, combined with polynomial time bounds for one trial of the $\Phi_k(p)$-method, shows that the whole procedure takes expected polynomial time.

4. Constructing Rings. Let m be a prime congruent to 1 modulo k, with a primitive root g. K_m denotes the cyclic field of degree k contained in the mth cyclotomic field,[†] and O_m denotes the ring of algebraic integers in K_m. Mathematically, R_m is just O_m reduced modulo N; this section tells how to implement it.

If ς is a primitive mth root of unity, then the *Gaussian period* of degree k is, by definition,

$$\eta = \sum_{x \in (\mathbf{Z}_m^*)^k} \varsigma^x.$$

Its conjugates are found by replacing ς by ς^g, ς^{g^2}, and so on in the above definition. This gives k quantities which we denote by $\eta_0 = \eta, \eta_1, \ldots, \eta_{k-1}$. K_m is constructed algebraically by adjoining a root of the irreducible polynomial

$$f(X) = (X - \eta_0)(X - \eta_1) \cdots (X - \eta_{k-1})$$

to the rational numbers. Its Galois group is generated by

$$\sigma : \eta_i \to \eta_{i+1 \bmod k}$$

(this depends on the generator g). Let Δ be the discriminant of f, modulo N; we may as well assume that $\gcd(\Delta, N)$ is 1 or N, for otherwise we get a factor of N.

[†]By the Kronecker-Weber theorem ([27, p. 341]), every abelian—hence every cyclic—field is contained in some cyclotomic field.

In the typical case, Δ is relatively prime to N, and then ([13, p. 27]) $R_m \cong \mathbf{Z}[X]/(f(X), N)$. Here we can use the power basis $\beta_i = X^i$ to implement R; the multiplication algorithms follow from polynomial algebra. The automorphism σ is implemented as a matrix: its coefficients t_{ij} have to satisfy

$$\sum_j t_{ij} \eta_0^j = \eta_1^i.$$

If N divides Δ, then we need to use the full ring of integers of K_m. This is a free $\mathbf{Z}$-module with basis $\eta_0, \eta_1, \ldots, \eta_{k-1}$ ([9, p. 217]), so we take $\beta_i = \eta_i$. In this case, σ is easy to implement (we just permute the basis elements cyclically). The multiplication tables can be computed by the following algorithm (see [26, Section 54]). Let $\bar{k} = (m-1)/k$, and

$$\iota(x) = (\text{index of } x \text{ in } \mathbf{Z}_m^*) \bmod k,$$
$$\alpha(x) = \begin{cases} \eta_{\iota(x)} & \text{if } (x, m) = 1, \\ \bar{k} \cdot \sum \eta_i & (\text{otherwise}) \end{cases}$$

(indices are taken with respect to the generator g). Then for $0 \le i, j < k$,

$$\eta_i \eta_j = \sum_{0 \le l < \bar{k}} \alpha(g^i + g^{kl+j}).$$

In practice, we only need to compute $\eta_0^2, \eta_0\eta_1, \ldots, \eta_0\eta_{k-1}$; the others are easily found using the Galois group. Finally, the multiplicative unit is

$$1 = -\eta_0 - \eta_1 - \cdots - \eta_{k-1}.$$

Even though K_m is defined as a subfield of a cyclotomic field, we would like to avoid using the larger field. For this reason, we suggest the following procedure.

First, find the multiplication tables in the period basis; this can be done by computing a table of indices (mod m) and using the algorithm above. Then, use these tables to express the quantities $1, \eta, \eta^2, \ldots, \eta^{k-1}$ as integral linear combinations of $\eta_0, \ldots, \eta_{k-1}$. Since the field discriminant of K_m is m^{k-1} ([18, p. 586]), the matrix of coefficients (t_{ij}) has the property that

$$\Delta = \text{disc}(f) = (\det(t_{ij}))^2 \cdot m^{k-1}$$

(apply [13, p. 64]). By applying Gaussian elimination mod N to (t_{ij}), we can see whether Δ is a unit mod N, since the determinant is the product of the pivots.

If Δ is a unit, then (since t_{ij} has been already made upper triangular) we can express η^k as a linear combination of smaller powers of η, giving the polynomial $f(X)$ mod N. To find σ, we find a relation

$$\sum_{i=0}^{k-1} x_i \eta^i = \eta'$$

and use this to express $(\eta')^2, (\eta')^3, \ldots$ in the power basis.

If Δ is not a unit, then (after seeing whether any of the pivots have a nontrivial gcd with N) we use the period basis.

Using classical algorithms for the arithmetic and Gaussian elimination for the linear algebra gives the operation counts below. In all cases, these estimates are good only up to constant factors.

COMPUTATION	POWER BASIS	PERIOD BASIS
1. multiplication table	–	$m \log^2 mk$
2. Gauss elimination	$k^3 \log^2 N$	–
3. find $f(X)$	$k \cdot \log^2 N$	–
4. find σ	$k^3 \log^2 N$	–
5. add/subtract in R	$k \cdot \log N$	$k \cdot \log N$
6. multiply/divide in R	$k^2 \log^2 N$	$k^3 \log^2 N$
7. apply σ	$k^3 \log^2 N$	$k \cdot \log N$

5. The Time for One Trial. Let N, the number to be factored, have at least two distinct prime factors p and q. We will say that a prime m is *useful* in separating p from q if it satisfies the following three conditions:[††]

1. m is congruent to 1 mod k,

2. p stays prime in O_m,

3. q splits completely in O_m.

In this section we analyze the running time of our procedure, assuming that such a prime has been found. We postpone the question of finding m until the next section.

We first need to estimate the coefficients of Ψ.

LEMMA 1. *Each coefficient of Ψ is bounded by 2^k in absolute value.*

Proof. The zeros of Ψ are roots of unity, at most k in number, so the ith coefficient is at most $\binom{k}{i}$. □

LEMMA 2. *The computation of x^E in R_m requires $O((\log E \cdot k^2 + k^3) \log^2 N)$ steps, if E is an integer.*

Proof. This is clear if we use the power basis. If R is implemented with the period basis, we use a special basis of powers of x, as follows. We first compute the matrix that represents multiplication by x, then apply this to powers of x to get x^2, $x^3, \ldots, x^k$ in the power basis. We use elimination (again!) to find a monic polynomial that x satisfies; all this requires $O(k^3)$ multiplications modulo N. Finally, we evaluate the power of x as usual, and express the result in the period basis. □

LEMMA 3. *The computation of $x^{\Psi(\tau)}$ requires $O(k^4 \log^2 N)$ steps.*

Proof. Since $\tau = \sigma^i$ with $\gcd(i, k) = 1$, we can first express $\Psi(\tau)$ as a polynomial in σ by rearranging coefficients. If we evaluate the symbolic power by a process

---

[††] We can also show that the complete splitting of q is not necessary, and also have a polynomial-time procedure that works as long as q does not stay prime. See Section 6.

similar to Horner's rule, we need $O(k)$ matrix-vector multiplications and $O(k)$ exponentiations by $O(k)$-bit numbers (by Lemma 1). The result follows from Lemma 2. □

LEMMA 4. *One trial of the Φ_k-method—that is, one execution of the outer loop, given a prime m as input—requires $O(m \log^2 m \cdot k + (k^5 + k^2 \log E) \log^2 N)$ steps.*

Proof. We compute t^E once, then compute $x^{\Psi(\tau)}$ at most k times. The result follows from Lemmas 2 and 3. □

The above results say that if m is small, then one trial of the algorithm will require polynomial time. How likely is this to produce a factorization? The answer is given next; this is the main technical result of the paper.

THEOREM 1. *Let N have at least two distinct prime factors p and q, and let m be useful in separating p from q. Let E be a multiple of $\Phi_k(p)$, such that $q - 1 \nmid E$. Let τ induce the Frobenius automorphism modulo p. Then, if t is a random element of R_m^*, with probability at least $1/2$, $t^{E \cdot \Psi(\tau)} \equiv 1 \pmod{p}$, $\not\equiv 1 \pmod{q}$. Furthermore, the time to compute this symbolic power of t is bounded by a polynomial in k, m, $\log E$, and $\log N$.*

Proof. If τ is the Frobenius mod p, then for any t in R_m^*, $t^{E \cdot \Psi(\tau)} \equiv 1 \bmod p$. We now have to show that the same relation is unlikely when p is replaced by q, or, what is the same thing, that the image of $t \to t^{E \cdot \Psi(\tau)} \pmod{q}$ consists of more than just the identity.

To prove this, let $Q_0, \ldots, Q_{k-1}$ be the prime ideals of O that divide q. It is known ([16, p. 70]) that the Galois group of K permutes them transitively; we assume that they are numbered so that $Q_i = Q_0^{\tau^i}$. The "residue notation"

$$x = (x_0, \ldots, x_{k-1})$$

indicates that $x \equiv x_i \pmod{Q_i}$. Here we can take the components x_i to be in $\mathbf{F}_q$.

Now, if $x - x_{k-1} \in Q_{k-1} = Q_0^{\tau^{k-1}}$ then (apply τ) $x^\tau - x_{k-1} \in Q_0$. This means that x^τ's first component is x_{k-1}; that is, the effect of τ is a right cyclic shift of the components x_i. Then, since Ψ is monic and has degree less than k,

$$(x, 1, \ldots, 1)^{\Psi(\tau)} = (\ldots, x, 1, \ldots, 1)$$

and

$$(x, 1, \ldots, 1)^{\Psi(\tau) \cdot E} = (\ldots, x^E, 1, \ldots, 1).$$

The set of such elements will have at least two elements when $q - 1 \nmid E$.

The statement about the running time follows from Lemma 4. □

6. The Nonsplitting Case. In this section we digress somewhat to show that the assumptions of the last section can be relaxed to include the case where q does not split completely. The results of this section are not required in any later sections. We will adhere to the notation of Section 5.

We will need some polynomial algebra first. Let k be a positive integer, and l a nontrivial divisor of k (that is, $l \mid k$ and $1 < l < k$). Let

$$\Psi_k(X) = \frac{X^k - 1}{\Phi_k(X)} = \Psi_k^{(0)}(X^l) + \Psi_k^{(1)}(X^l) \cdot X + \cdots + \Psi_k^{(l-1)}(X^l) \cdot X^{l-1}$$

(here we are just grouping terms according to their degree modulo l). It will be useful to have an explicit form for $\Psi_k^{(0)}$. Let ς denote a primitive lth root of unity;

then, since

$$\Psi_k^{(0)}(X^l) = \frac{1}{l}[\Psi_k(X) + \Psi_k(\varsigma X) + \cdots + \Psi_k(\varsigma^{l-1}X)],$$

we see that

(2) $$\Psi_k^{(0)}(X) = \frac{1}{l}[\Psi_k(X^{1/l}) + \Psi_k(\varsigma X^{1/l}) + \cdots + \Psi_k(\varsigma^{l-1}X^{1/l})].$$

LEMMA 5. *If ε is a root of unity of order t for some $t \mid k/l$, $t \neq k/l$, then $\Psi_k^{(0)}(\varepsilon) = 0$.*

Proof. Let ς denote a primitive lth root of unity. Then for any i, $0 \leq i < l$,

$$(\varsigma^i \varepsilon^{1/l})^{tl} = \varsigma^{itl}\varepsilon^t = 1.$$

Now, since $t \mid k/l$ and $t \neq k/l$, $tl \mid k$ and $tl \neq k$. Hence, $\varsigma^i \varepsilon^{1/l}$ is a root of Ψ_k. (Ψ_k is the product of all cyclotomic polynomials of order properly dividing k.) By (2), then,

$$\Psi_k^{(0)}(\varepsilon) = 0. \quad \square$$

LEMMA 6. *The polynomials $\Psi_k^{(0)}(X)$ and $\Phi_{k/l}(X)$ are relatively prime.*

Proof. Recall that

$$X^{k/l} - 1 = \prod_{t \mid k/l} \Phi_t(X).$$

Since $\Psi_k^{(0)}$ has degree less than k/l, it cannot contain all $\Phi_t(X)$'s as factors. However, by Lemma 5, it must contain factors of the form $\Phi_t(X)$ when $t \mid k/l$, $t \neq k/l$, and therefore $\Phi_{k/l}$ must be the one omitted. The result follows. $\square$

Let $R(f,g)$ denote the resultant of the polynomials f and g. For k and l as above, define

(3) $$R_{k,l} = R(\Psi_k^{(0)}(X), \Phi_{k/l}(X)).$$

From standard properties of resultants ([12, p. 135]) and Lemma 6 it follows that $R_{k,l}$ is a nonzero integer contained in the ideal generated by $\Psi_k^{(0)}$ and $\Phi_{k/l}$ in $\mathbf{Z}[X]$. We also have the alternative expression

$$R_{k,l} = R(\Psi_k^{(0)}(X^i), \Phi_{k/l}(X))$$

whenever i is relatively prime to k/l. This follows from the lemma below.

LEMMA 7. *Let $f(X) \in \mathbf{Z}[X]$, Φ_a the cyclotomic polynomial of order a. Then, if i is relatively prime to a,*

$$R(f(X), \Phi_a(X)) = R(f(X^i), \Phi_a(X)).$$

Proof. If c is the leading coefficient of f, and ϕ denotes the Euler phi-function, by a known formula ([12, p. 137])

$$R(f(X), \Phi_a) = c^{\phi(a)} \prod (\alpha - \beta);$$

here the product is taken over all roots α of f and β of Φ_a (taken with appropriate multiplicities). Then

$$R(f(X^i), \Phi_a) = c^{\phi(a)} \prod (\alpha' - \beta),$$

where α' denotes a root of $f(X^i)$. Since each α' has to be of the form $\varsigma^j \alpha^{1/i}$, where ς is a primitive ith root of unity, we can group the factors corresponding to a given α together to get

$$(\alpha^{1/i} - \beta)(\varsigma \alpha^{1/i} - \beta) \cdots (\varsigma^{i-1}\alpha^{1/i} - \beta) = \alpha - \beta^i.$$

But $\beta \to \beta^i$ just permutes the roots of Φ_a, and the result follows. $\square$

Resultants can be computed efficiently by the Euclidean algorithm ([3, p. 160]); however, for the purposes of this discussion we can think of them as precomputed.

THEOREM 2. *Let q be a prime dividing N that is unramified in $K_m/\mathbf{Q}$, and splits into l pieces, where $1 < l < k$. Define $R_{k,l}$ as in (3). Assume that*

$$\Phi_{k/l}(q) \nmid R_{k,l}E.$$

Let τ generate the Galois group of $K_m/\mathbf{Q}$. Then there is some $x \in R_m^*$ for which

$$x^{\Psi_k(\tau)E} \not\equiv 1 \pmod{q}.$$

Proof. Let $Q_0 \cdots Q_{l-1}$ denote the prime ideals dividing q; then for all i, $0 \le i < l$, $R_m/Q_i \cong \mathbf{F}_{q^{k/l}}$. Since τ permutes the prime ideals dividing q cyclically, τ_l induces an automorphism of R_m/q, given by

$$x^{\tau^l} \equiv x^{q^i} \pmod{q}$$

for some i relatively prime to k/l. Therefore (using residue notation),

$$(x, 1, \ldots, 1)^{\Psi_k(\tau)} \equiv (x^{\Psi_k^{(0)}(q^i)}, \ldots)$$

and so

$$x^{\Psi_k(\tau)E} \equiv x^{\Psi_k^{(0)}(q^i)E} \pmod{Q_0}.$$

This cannot be identically 1 for $x \in (R_m/Q_0)^*$ unless $q^{k/l} - 1 \mid \Psi_k^{(0)}(q^i)E$. But this last condition would imply

$$\Phi_{k/l}(q) | \Psi_k^{(0)}(q^i)E.$$

Recall the assumption that

$$\Phi_{k/l}(q) \nmid R_{k,l}E.$$

But by properties of resultants, we have polynomials α_i and β_i with integer coefficients such that

$$R_{k,l}E = \alpha_i(q)\Psi_k^{(0)}(q^i)E + \beta_i(q)\Phi_{k/l}(q)E.$$

The right side of the above equation is clearly divisible by $\Phi_{k/l}(q)$, but the left side is not, a contradiction. $\square$

We now indicate how one could prove an analog of Theorem 1. Assume that N, the number to be factored, has two distinct prime factors p and q. Say that an auxiliary prime $m \equiv 1 \pmod{k}$ is *good* if:

1. p stays prime in O_m
2. q does not stay prime in O_m.

Such a prime can be used in a recursive factorization algorithm outlined below. This procedure runs in polynomial time and will split N with high probability

provided that m is good and $\Phi_k(p) \mid E$; a proof of this fact is left to the reader. The base case is $k = 1$, which is the $p - 1$ method.

THE CYCLOTOMIC FACTORING ALGORITHM

Cyclo (k, E):

For each $l \mid k$, $l \neq 1$:

Compute the polynomial $\Psi_k^{(0)}(X)$.

Compute $R_{k,l} = R(\Psi_k^{(0)}, \Phi_{k/l})$.

Run Cyclo $(k/l, R_{k,l}E)$ to ensure that $\Phi_{k/l}(q) \nmid R_{k,l}E$.

For each i, $1 \leq i \leq k$, relatively prime to k:

Choose a random $x \in R_m$.

Try to split N with $\gcd(x^{\Psi_k(\tau)E} - 1, N)$.

7. Higher Reciprocity. The use of generalized reciprocity laws in algorithms may strike the reader as a black art. Since these laws are essential for an understanding of our algorithm, we present some background; the viewpoint is that of [28].

We are concerned with the following situation. Let $f(X)$ be an irreducible monic polynomial with integral coefficients that is *normal* in the sense of field theory. That is, adjoining one root of f to the rationals creates a field K that contains all roots of f. To properly analyze our algorithm, we will have to answer the following kind of question.

Let p be a prime number. How will f factor modulo p? Furthermore, what happens to the Galois group mod p?

One can give precise answers to these questions in the case where f is *abelian*, that is, K has a commutative Galois group.

The splitting information is traditionally encoded in the following fashion. For most[†††] primes p, the defining polynomial $f(X)$ will factor modulo p into r polynomials $f_1, \ldots, f_r$ of the same degree. Then (applying [13, p. 18]) there is a unique element σ of the Galois group for which

$$X^p \equiv X^\sigma \pmod{f_i}$$

for all i. This is called the *Artin symbol* of p, and written

$$\sigma = (p \mid K/\mathbf{Q}).$$

For our algorithm, the most important fact is this: the order of p's Artin symbol is equal to the degree of one (hence every) factor f_i.

Clearly, then, to find out how polynomials split, we must compute Artin symbols. How do we do this? First consider the classical quadratic reciprocity law, which states that for distinct odd primes p and m, $X^2 - p$ splits mod m in the same way as $X^2 - m^*$ splits mod p (here m^* is $\pm m$, whichever is congruent to 1 mod 4). Notice that we are not strictly interchanging the roles of p and m; we construct a new equation of the same degree (somehow using m), and ask how it splits mod p. For arbitrary $k > 1$, the generalization is given below in Lemma 10.

[†††]The exceptions are the primes dividing the discriminant of f (which are finite in number).

LEMMA 8. *Let m be a prime and let the finite field $\mathbf{Z}_m$ contain a kth root of unity; that is, $m \equiv 1 \pmod{k}$. Let G be a multiplicative group with*

$$(\mathbf{Z}_m^*)^k \subset G \subset \mathbf{Z}_m^*$$

and let K_G denote the field $\mathbf{Z}_m(G^{1/k})$ (that is, we adjoin a kth root of every element of G to $\mathbf{Z}_m$). Then the degree of the field extension $K_G/\mathbf{Z}_m$ is equal to the index of $(\mathbf{Z}_m^*)^k$ in G.

Proof. See the discussion of Kummer theory in [12]. $\square$

LEMMA 9. *The order of p in $\mathbf{Z}_m^*/(\mathbf{Z}_m^*)^k$ is the degree of any irreducible factor of $X^k - p$ modulo m.*

Proof. Apply Lemma 8 when G is the group generated by p. It is true, then, that

$$\mathbf{Z}_m(G^{1/k}) = \mathbf{Z}_m(\sqrt[k]{p}).$$

But $\mathbf{Z}_m(\sqrt[k]{k})$ is the splitting field of $X^k - p$ and the index of G is just the order of p. Let

$$X^k - p = f_1(X) \cdots f_r(X)$$

be the complete factorization over $\mathbf{Z}_m$. If we adjoin any root of an f_i to $\mathbf{Z}_m$, we get a field containing the roots of all f_i's, since they only differ by a kth root of unity. We conclude that all the f_i's have the same degree, which is equal to the degree of the splitting field and hence equal to the order of p modulo the group of kth powers. $\square$

LEMMA 10. *Let m be prime, congruent to 1 mod k, and let $f_m(X)$ be the period polynomial of degree k, defined in Section 4. Then, if $p \nmid \operatorname{disc}(f_m)$, $X^k - p$ splits mod m in the same way as $f_m(X)$ splits mod p.*

Proof. Let η (the period) be a root of $f_m(X)$. Since p does not divide the discriminant of f_m, p splits in exactly the same way as f_m splits mod p, so we have just to consider the splitting of the prime ideal (p) in going from $\mathbf{Q}$ to $\mathbf{Q}(\eta)$.

To do this, consider $K = \mathbf{Q}(\eta)$ as a subfield of the cyclotomic field $\mathbf{Q}(\varsigma)$ (ς is a primitive mth root of unity). Then there is a unique subgroup G of $\mathbf{Z}_m^*$ with the following property: $\mathbf{Z}_m^*/G$ is isomorphic to $\operatorname{Gal}(K/\mathbf{Q})$, and the isomorphism is induced by the map $p \to (p \mid K/\mathbf{Q})$ (use Theorem 2 of [27, p. 338]).

By Galois theory, the index $(\mathbf{Z}_m^* : G)$ has to be k, the degree of K, so that G is just $(\mathbf{Z}_m^*)^k$. As remarked above, the order of $(p \mid K/\mathbf{Q})$ has to be the residue degree of any of its prime divisors, otherwise the Artin symbol would not be unique.

This gives the result, by Lemma 9. $\square$

In the next section, we will have to use a more general version of the above theory. First, we have to consider not just the splitting of polynomials, but also the splitting of prime ideals. Second, we consider not just extensions of the rationals, but any relatively-abelian extension of algebraic number fields. We need the first generalization because in the "exceptional" case, we cannot use a polynomial ring, but have to consider the full ring of integers of K_m. To justify the second generalization, consider that we are interested in the splitting of polynomials like

(4) $$X^k - p \pmod{m}.$$

At first glance, the theory will not apply, since this equation is not even normal. However, we can get around this problem by considering an extension field over which (4) is abelian. Then, if we can find a prime ideal M of this field with norm m, the splitting of (4) mod m will translate into the splitting of M (since the equation has rational integral coefficients).

8. A Bound for Useful Primes. This section proves a technical result (Theorem 3, below) which gives a polynomial bound on the least m that makes the Φ_k-method work. The result is not hard to state, but the proof requires some technicalities of algebraic number theory, in particular the notion of Artin symbols for relatively-abelian extensions ([27, p. 338]).

We make the following assumptions: m, p, and q are distinct primes, and k is a divisor of $m-1$, relatively prime to pq. K_m is the unique field of degree k contained in $\mathbf{Q}(\varsigma_m)$; its ring of integers is denoted O_m. $f_m(X)$ is the period polynomial as defined in Section 4, with discriminant Δ.

LEMMA 11. *p stays prime in O_m and q splits completely in O_m if and only if p generates $\mathbf{Z}_m^*/(\mathbf{Z}_m^*)^k$ and q is a kth power modulo m. Furthermore, if $\gcd(pq, \Delta) = 1$, this happens if and only if f_m is irreducible mod p and splits completely mod q.*

Proof. This follows from the proof of Lemma 10. □

Where p and q lie in the group $\mathbf{Z}_m^*/(\mathbf{Z}_m^*)^k$ depends on the splitting of $X^k - p$ and $X^k - q$ modulo m. To discuss this, we introduce two auxiliary fields:

$$L = \mathbf{Q}(\omega, \sqrt[k]{q}), \qquad L' = \mathbf{Q}(\omega, \sqrt[k]{q}, \sqrt[k']{p}).$$

(Here, ω is a primitive kth root of unity, and k' is the maximal squarefree divisor of k.) Then L' is a cyclic extension of L; let A' and A denote the respective rings of integers of these two fields.

LEMMA 12. *Let M be a prime ideal of A of degree 1 whose Artin symbol $(M \mid L'/L)$ has order k'. Then the rational prime $m = \text{Norm}(M)$ has the following properties: $m \equiv 1 \pmod{k}$, p generates $\mathbf{Z}_m^*/(\mathbf{Z}_m^*)^k$, and $q \in (\mathbf{Z}_m^*)^k$.*

Proof. First, since M has degree 1 in L, m splits completely in any subfield of L. Taking the subfield to be $\mathbf{Q}(\omega)$, we see that $m \equiv 1 \pmod{k}$.

Next, since the order of its Artin symbol is as large as possible, M stays prime in going from L to L'. This implies that $X^{k'} \equiv p \pmod{M}$ cannot be solved in A. Since M has degree 1, $A/MA \cong \mathbf{F}_m$, so $X^{k'} - p$ is irreducible mod m. Therefore, p is not an rth power mod m for any prime r dividing k', so p has to generate $\mathbf{Z}_m^*/(\mathbf{Z}_m^*)^k$.

Finally, M has to lie above some prime ideal of norm m in $\mathbf{Q}(\omega)$. This splits completely in going to L, so (by the above argument) $X^k - q$ has to factor completely modulo m, meaning that $q \in (\mathbf{Z}_m^*)^k$. □

LEMMA 13. *Let $L \subset L'$ be an abelian extension of number fields, and let σ be in $\text{Gal}(L'/L)$. Then there is a prime ideal M of L with residue degree 1 such that the Artin symbol $(M \mid L'/L) = \sigma$. If the Dedekind zeta function of L' satisfies the Riemann hypothesis, then there is such an M with $\text{Norm}(M) = O(\log^2 |\Delta'|)$, where Δ' is the discriminant of L'.*

Proof. The existence of a small prime ideal with the right Artin symbol is from [10, Corollary 1.2], specialized to relatively-abelian extensions. One can further take this ideal to have residue degree 1, as remarked in the introduction to [11]. □

LEMMA 14. *If k, p, q, and L' are as defined above, then*

$$\log |\operatorname{disc} L'| \le 3k^3 \log pqk.$$

Proof. It is known that the discriminant of $X^n - a$ is $\pm n^n a^{n-1}$ (e.g., [6, p. 91]). Therefore, the discriminants of $\mathbf{Q}(\omega)$, $\mathbf{Q}(\sqrt[k]{q})$, and $\mathbf{Q}(\sqrt[k]{p})$ divide k^k, $(kq)^k$, and $(kp)^k$, respectively. By an estimate for the discriminants of composed fields ([21, Lemma 7]).

$$|\operatorname{disc} \mathbf{Q}(\omega, \sqrt[k]{q}, \sqrt[k]{p})| \le k^{3k^3} \cdot (pq)^{k^3}. \quad \square$$

THEOREM 3. *Let k be a positive integer. The least prime $m \equiv 1$ modulo k for which p stays prime in O_m and q splits completely is, assuming ERH, $O(k^6 \log^2(pqk))$.*

Proof. Let σ be a generator for $\operatorname{Gal}(L'/L)$, and apply Lemmas 12, 13, and 14. □

9. Conclusion. Our results above can be summed up in the following two theorems.

THEOREM 4. *Let k be a positive integer. Let Φ_k be the kth cyclotomic polynomial. Then, assuming ERH, there is a random polynomial-time algorithm that takes as input a multiple E of $\Phi_k(p)$ and splits any multiple of p.*

Proof. The algorithm proceeds as follows. First check that N is not a prime power. Run the Φ_k-method, using primes m of the form $jk + 1$, $j = 1, 2, \ldots$, in order up to the bound indicated by Theorem 3. In parallel, run the Φ_1-method (that is, the $p - 1$ method); note that here the auxiliary prime is not necessary since we can do everything in the field of rational numbers. By Theorems 1 and 2, with probability $1/2$ we either factor using one of the small auxiliary primes, or we factor using the $p - 1$ method. The whole process takes polynomial time. □

THEOREM 5. *Let A and B be positive integers. Call a number N vulnerable if it has a prime factor p with the following property: for some $k \le A$, $\Phi_k(p)$ consists of primes less than B. Then, assuming ERH, there is a random polynomial-time algorithm to split all vulnerable numbers.*

Proof. Since $\Phi_k(p) \le N^A - 1$, use $E = \prod_{p \le B} p^{\lceil A \log_p N \rceil}$ as input to all the Φ_k-methods. By the prime number theorem, $\log E \le A \cdot \pi(B) \log N$, giving a polynomial-time bound. □

Part of our interest in these methods comes from the following question: "What fraction of numbers can we expect to be able to factor?" The answer certainly hinges on what the density of vulnerable numbers is, as a function of A and B.

Algebraically, all the cyclotomic factoring methods boil down to the following idea. Let p be a prime factor that we want to remove, and consider the field extension $\mathbf{F}_{p^k}/\mathbf{F}_p$. This has a cyclic Galois group G, generated by (say) σ. Then $\mathbf{F}_{p^k}^*$ is a module over the group ring $\mathbf{Z}[G]$, which we know is isomorphic to $\mathbf{Z}[X]/(X^k - 1)$. If $f(X) \mid X^k - 1$, the set of multiples of $f(\sigma)$ is a submodule A_f of $\mathbf{F}_{p^k}^*$. We get

interesting results when σ is the Frobenius mod p, and we can annihilate A_f *without* knowing p. This explains the usefulness of cyclotomic polynomials in factoring: they are intimately related to the structure of the group ring.

In our algorithm, we try to annihilate A_{Φ_k}: to do this, we must construct a global field with an automorphism that does what we want mod p. An interesting question is whether this expensive "globalization" of the problem is necessary; after all, we are really only interested in what happens modulo p and q.

We can give a partial answer using the results of [24], which we paraphrase as follows. Let r and s be two distinct primes, and let L/K be a cyclic field extension of degree $k = rs$, with Galois group generated by σ. Then K has two subfields K_r and K_s, of relative degrees r and s, respectively; let N be the norm from L to K_s and T the trace (or any other symmetric function) from L to K_s. Then, if t is chosen so that $N(t) = 1$,

$$x = t^{\Phi_k(\sigma)} \in K_r,$$

and therefore $T(x) \in K$.

The nice thing about this result is that we do not need to know σ at all; if the base field K is $\mathbf{F}_p$, it suffices to know $\Phi_k(p)$ (see [2] for the details). This gives a simple "local" factoring method, provided that k has at most two distinct prime factors; we do not yet know how to make it work for general k.

Finally, we would be amiss not to mention a closely related primality test, given by Lenstra [14] and relying on the following observation. Let p be a prime, and suppose we know the complete factorization of some divisor m of $p^k - 1$. Then if $m > \sqrt{p}$, we can easily prove p prime. Taking $m = \Phi_k(p)$, we get a "cyclotomic" prime test; there is also the nice possibility of combining partial factorizations of various $\Phi(p)$'s in constructing a primality proof. It would be of interest to see if partial information about several $\Phi(p)$'s could similarly be used in factoring.

Acknowledgments. One of us (Bach) learned about the $\Phi(p)$-methods from Martin Hellman, and has been fascinated ever since. We would also like to thank Hendrik Lenstra for suggesting the use of the period basis, and Gary Shute for suggesting an algorithm for exponentiation using periods.

Appendix: Three Examples. In this appendix we present three examples in detail, for the Φ_2, Φ_4, and Φ_6 methods.

Example 1: $k = 2$. Recall that we are going to factor N, with two distinct prime divisors p and q. We present algebraic details for the Φ_2 method, using the period basis.

First we find a prime $m \equiv 1 \pmod{k}$; we hope that p generates $\mathbf{Z}_m^*/(\mathbf{Z}_m^*)^k$ (i.e., p is a quadratic residue mod m) and q is a kth power $\pmod{m}$ (i.e., q is a quadratic nonresidue mod m). We therefore hope that the Jacobi symbol

$$(pq \mid m) = -1.$$

Next we construct the periods; they are

$$\eta_0 = \sum_{1 < k < m/2} \varsigma^{k^2}$$

and

$$\eta_1 = \sum_{1 < k < m/2} \varsigma^{rk^2} \qquad (r \text{ a nonresidue}).$$

The period equation is

$$f(X) = (X - \eta_0)(X - \eta_1) = X^2 - (\eta_0 + \eta_1)X + \eta_0\eta_1 = 0.$$

Here, $\eta_0 + \eta_1 = -1$ (this is easy) and $\eta_0\eta_1 = [1 - (-1 \mid m)\,m]/4$ (this is nontrivial; it is a result of Gauss ([7, Section 355])). From this we see that

$$\eta_{0,1} = \frac{-1 \pm \sqrt{m^*}}{2},$$

where $m^* = (-1 \mid m)\,m$. Next we need a generator for the Galois group; this is a 2×2 matrix σ such that

$$\begin{pmatrix} 1 \\ \eta_1 \end{pmatrix} = \sigma \begin{pmatrix} 1 \\ \eta_0 \end{pmatrix}.$$

We find

$$\sigma = \begin{pmatrix} 1 & 0 \\ -1 & -1 \end{pmatrix}.$$

Finally, $\Phi_2(X) = X + 1$ and $\Psi_2(X) = X - 1$. This concludes the "precomputation" phase of the algorithm.

For one trial, we choose at random $x = a_0 + a_1\eta_0$ and compute

$$x^{\Psi_2(\sigma)} = x^{\sigma-1} = \frac{a_0 - a_1 - a_1\eta_0}{a_0 + a_1\eta_0} = y.$$

If we are successful, then $\gcd(y^E - 1, N)$ splits N. It will be seen that this reduces essentially to the old $p + 1$ method.

Example 2: $k = 4$. Here we show how to construct the Galois group generator matrix σ for the case $k = 4$, $m = 13$. Let ς denote a primitive 13th root of unity. The periods are

$$\eta_0 = \varsigma + \varsigma^3 + \varsigma^9,$$
$$\eta_1 = \varsigma^2 + \varsigma^6 + \varsigma^5,$$
$$\eta_2 = \varsigma^4 + \varsigma^{12} + \varsigma^{10},$$
$$\eta_3 = \varsigma^8 + \varsigma^{11} + \varsigma^7.$$

From this we find the irreducible polynomial satisfied by the periods; it is

$$f(X) = (X - \eta_0)(X - \eta_1)(X - \eta_2)(X - \eta_3) = X^4 + X^3 + 2X^2 - 4X + 3.$$

By using $f(\eta_i) = 0$, we find

$$\eta_0^2 = -2\eta_3 - 2\eta_0 - \eta_1 - 2 = \eta_1 + 2\eta_2$$

and

$$\eta_0^3 = 3\eta_3 + \eta_0 + 3\eta_1 + 6 = -5\eta_0 - 3\eta_1 - 6\eta_2 - 3\eta_3.$$

This gives us the linear system

$$1 = -\eta_0 - \eta_1 - \eta_2 - \eta_3,$$
$$\eta_0 = \eta_0,$$
$$\eta_0^2 = \eta_1 + 2\eta_2,$$
$$\eta_0^3 = -5\eta_0 - 3\eta_1 - 6\eta_2 - 3\eta_3,$$

which can be easily solved to get

$$(5) \qquad\qquad \eta_1 = \frac{1}{3}(2\eta_0^3 + 3\eta_0^2 + 4\eta_0 - 6).$$

Thus, we can fill in the first two lines in the matrix equation

$$\begin{pmatrix} 1 \\ \eta_1 \\ \eta_1^2 \\ \eta_1^3 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ -2 & 4/3 & 1 & 2/3 \\ & & & \\ & & & \end{pmatrix} \begin{pmatrix} 1 \\ \eta_0 \\ \eta_0^2 \\ \eta_0^3 \end{pmatrix}.$$

To fill in the rest of the lines, we just use (5) and the fact that $f(\eta_i) = 0$ to get

$$\eta_1^2 = \left(\frac{1}{3}(2\eta_0^3 + 3\eta_0^2 + 4\eta_0 - 6)\right)^2 = -\eta_0^3 - 2\eta_0^2 - 4\eta_0 + 1$$

and

$$\eta_1^3 = \left(\frac{1}{3}(2\eta_0^3 + 3\eta_0^2 + 4\eta_0 - 6)\right)^3 = -\frac{1}{3}\eta_0^3 + \eta_0^2 + \frac{7}{3}\eta_0 + 7.$$

Thus, we finally obtain

$$\begin{pmatrix} 1 \\ \eta_1 \\ \eta_1^2 \\ \eta_1^3 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ -2 & 4/3 & 1 & 2/3 \\ 1 & -4 & -2 & -1 \\ 7 & 7/3 & 1 & -1/3 \end{pmatrix} \begin{pmatrix} 1 \\ \eta_0 \\ \eta_0^2 \\ \eta_0^3 \end{pmatrix}$$

(notice that the denominators are at worst 3; this is because the discriminant of f is $3^2 \cdot 13^3$).

Example 3: $k = 6$. We factor $N = 1142624627800367$ using the Φ_6-method. The reader is encouraged to follow along using a computer algebra system such as MACSYMA.

Choose $m = 7$; then K_7 is a cyclotomic field, generated by ς, a root of

$$f(X) = X^6 + X^5 + X^4 + X^3 + X^2 + X + 1.$$

Since 3 is a generator for $\mathbf{Z}_6^*$, the Galois group of K_7 is generated by $\varsigma \rightarrow \varsigma^3$. This allows us to easily work out the matrix for σ relative to the power basis $(1, \varsigma, \varsigma^2, \dots, \varsigma^5\}$ (it is almost a permutation matrix). Since

$$\Phi_6(X) = X^2 - X + 1,$$

we find

$$\Psi_6(X) = X^4 + X^3 - X - 1.$$

Let $t = 2 + 3\varsigma$; then

$$t^{\Psi_6(\sigma)} = 56761050625072\varsigma^5 - 579949865082260\varsigma^4 - 360309277880893\varsigma^3$$
$$+ 841544272310854\varsigma^2 + 693471966332404\varsigma - 594757095680104$$

and

$$[t^{\Psi_6(\sigma)}]^{300!} = -236435130228200\varsigma^5 + 26173735275454\varsigma^4 + 497751375820655\varsigma^3$$
$$+ 107621379503601\varsigma^2 - 79744299524943\varsigma + 333155860818985.$$

Finally, $\gcd(-236435130228200, N) = 149861$, so the factorization is

$$N = 149861 \cdot 7624562947$$

(7624562947 is prime). We were successful because

$$\Phi_6(149861) = 3 \cdot 7^2 \cdot 13 \cdot 19 \cdot 37 \cdot 73 \cdot 229.$$

Computer Sciences Department
University of Wisconsin
Madison, Wisconsin 53706
E-mail: bach@cs.wisc.edu

Department of Mathematics and Computer Science
Dartmouth College
Hanover, New Hampshire 03755
E-mail: shallit@dartmouth.edu

1. E. BACH, G. MILLER & J. SHALLIT. "Sums of divisors, perfect numbers, and factoring," *SIAM J. Comput.*, v. 15, 1986, pp. 1143–1154.

2. E. BACH & J. SHALLIT, *A Class of Functions Equivalent to Factoring*, Technical Report 84–008, Department of Computer Science, University of Chicago, 1984.

3. E. R. BERLEKAMP, *Algebraic Coding Theory*, McGraw-Hill, New York, 1968.

4. J. BRILLHART, D. H. LEHMER, J. SELFRIDGE, B. TUCKERMAN & S. S. WAGSTAFF, JR., *Factorizations of* $b^n - 1$, $b = 2, 3, 5, 6, 7, 10, 11, 12$ *Up to High Powers*, Contemp. Math., vol. 22, Amer. Math. Soc., Providence, R. I., 1983.

5. J. D. DIXON, "Factorization and primality tests," *Amer. Math. Monthly*, v. 91, 1984, pp. 333–352.

6. J. W. CASSELS & A. FRÖHLICH, *Algebraic Number Theory*, Academic Press, London, 1976.

7. C. F. GAUSS, *Disquisitiones Arithmeticae*, Springer-Verlag, Berlin, 1986.

8. R. K. GUY, *How to Factor a Number*, Proc. 5th Manitoba Conf. on Numerical Mathematics, 1975, pp. 49–89.

9. D. HILBERT, *Die Theorie der algebraischen Zahlkörper*, Jahresbericht der Deutschen Mathematikervereinigung, vol. 4, 1897, pp. 175–546. [This has been reprinted in vol. 1 of Hilbert's collected works by Chelsea, 1981.]

10. J. C. LAGARIAS & A. M. ODLYZKO, "Effective versions of the Chebotarev density theorem," in *Algebraic Number Fields* (A. Fröhlich, ed.), pp. 409–464, Academic Press, London, 1977.

11. J. C. LAGARIAS, H. L. MONTGOMERY & A. M. ODLYZKO, "A bound for the least prime ideal in the Chebotarev density theorem," *Invent. Math.* v. 54, 1979, pp. 271–296.

12. S. LANG, *Algebra*, Addison-Wesley, Reading, Mass., 1971.

13. S. LANG, *Algebraic Number Theory*, Addison-Wesley, Reading, Mass., 1970.

14. H. W. LENSTRA, JR., Primality Testing Algorithms [after Adleman, Rumley, and Williams], *Séminaire Bourbaki*, in: Lecture Notes in Math., vol. 901, pp. 576–01 to 576–15, Springer, Berlin, 1981.

15. H. W. LENSTRA, JR., "Factoring integers with elliptic curves," *Ann. of Math.*, v. 126, 1987, pp. 649–673.

16. D. A. MARCUS, *Number Fields*, Springer, New York, 1977.

17. G. MILLER, "Riemann's hypothesis and tests for primality," *J. Comput. System Sci.*, v. 13, 1976, pp. 300–317.

18. E. NETTO, "Ueber die Factorenzerlegung der Discriminanten algebraischer Gleichungen," *Math. Ann.*, v. 24, 1884, pp. 579–587.

19. J. M. POLLARD, "Theorems on factorization and primality testing," *Proc. Cambridge Philos. Soc.*, v. 76, 1974, pp. 521–528.

20. C. P. SCHNORR & H. W. LENSTRA, JR., "A Monte-Carlo factoring algorithm with linear storage," *Math. Comp.*, v. 43, 1984, pp. 289–311.

21. H. STARK, "Some effective cases of the Brauer-Siegel theorem," *Invent. Math.*, v. 23, 1974, pp. 135–172.

22. H. C. WILLIAMS, "A generalization of Lehmer's functions," *Acta Arith.*, v. 24, 1976, pp. 315–341.

23. H. C. WILLIAMS, "A $p + 1$ method of factoring," *Math. Comp.*, v. 39, 1982, pp. 225–234.

24. H. C. WILLIAMS & J. S. JUDD, "Determination of the primality of N by using factors of $N \pm 1$," *Math. Comp.*, v. 30, 1976, pp. 157–172.

25. H. C. WILLIAMS & J. S. JUDD, "Some algorithms for prime testing using generalized Lehmer functions," *Math. Comp.*, v. 30, 1976, pp. 867–886.

26. B. L. VAN DER WAERDEN, *Modern Algebra*, Ungar, New York, 1950. [The relevant material on cyclotomic periods does not appear in the second English edition.]

27. L. WASHINGTON, *Cyclotomic Fields*, Springer, New York, 1982.

28. B. F. WYMAN, "What is a reciprocity law?," *Amer. Math. Monthly*, v. 79, 1972, pp. 571–586.