# An Implementation of the Number Field Sieve 

Marije Elkenbracht-Huizing

## CONTENTS

1. Introduction
2. Description of the NFS
3. Outline of the Implementation
4. Free Relations
5. Choice of the Polynomials
6. The Sieving
7. The Filtering
8. The Block Lanczos Method
9. Extracting the Square Root
10. Experimental Results

Acknowledgements
References

AMS Subject Classification (1991): 11Y05, 11Y40
Keywords: number field sieve, factorization.
This research is funded by the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (Netherlands Organization for Scientific Research, NWO) through the Stichting Mathematisch Centrum (SMC), under grant number 611-307-022.

The Number Field Sieve (NFS) is the asymptotically fastest known factoring algorithm for large integers. This article describes an implementation of the NFS, including the choice of two quadratic polynomials, both classical sieving and a special form of lattice sieving (line sieving), the block Lanczos method and a new square root algorithm. Finally some data on factorizations obtained with this implementation are listed, including the record factorization of $12^{151}-1$.

## 1. INTRODUCTION

The Number Field Sieve (NFS), introduced in 1988 [Pollard 1993a], is the asymptotically fastest known algorithm for factoring integers. Two forms of the NFS have been considered: the Special NFS, or SNFS, tailored especially to integers of the form $n=c_{1} r^{t}+c_{2} s^{u}$, and the General NFS, or GNFS, applicable to arbitrary numbers. The NFS factors integers $n$ in heuristic time

$$
\exp \left((c+o(1))(\log n)^{1 / 3}(\log \log n)^{2 / 3}\right)
$$

as $n \rightarrow \infty$, where $c=\left(\frac{32}{9}\right)^{1 / 3} \approx 1.5$ for the SNFS and $c=\left(\frac{64}{9}\right)^{1 / 3} \approx 1.9$ for the GNFS [Buhler et al. 1993]. These expressions should be compared with the time

$$
\exp \left((1+o(1))(\log n)^{1 / 2}(\log \log n)^{1 / 2}\right)
$$

taken by the Multiple Polynomial Quadratic Sieve, or MPQS [Pomerance 1985], still the best generalpurpose factoring algorithm for integers with less than approximately 105 digits.

We describe here several experiments carried out with an implementation of the NFS written by J. Buhler, R. M. Elkenbracht-Huizing, P. L. Montgomery, R. Robson and R. Ruby. It has been used, among others, for the record SNFS factorization
of $\left(12^{151}-1\right) / 11$, a number of 162 decimal digits, and a GNFS factorization of a 107-digit cofactor of $6^{223}+1$. We start with a description of the NFS and an outline of the implementation, then discuss in more detail several aspects of the implementation, and finally state the results of the factorization experiments. Detailed descriptions of the NFS can be found in [Lenstra et al. 1993b; Buhler et al. 1993].

## 2. DESCRIPTION OF THE NFS

Let $n$ be the odd number to be factored. It is easy to check whether $n$ is a prime number or a prime power [Lenstra et al. 1993c, § 2.5], and we assume that it is neither. Like MPQS, the NFS tries to find a solution of the equation $v^{2} \equiv w^{2} \bmod n$. For at least half of the pairs $(v \bmod n, w \bmod n)$ with $v^{2} \equiv w^{2} \bmod n$ and $v$ and $w$ relatively prime to $n$, the greatest common divisor of $n$ and $v-w$ gives a nontrivial factor of $n$.

To construct $v$ and $w$ we first choose two polynomials

$$
\begin{aligned}
& f_{1}(x)=c_{1, d_{1}} x^{d_{1}}+c_{1, d_{1}-1} x^{d_{1}-1}+\cdots+c_{1,0} \\
& f_{2}(x)=c_{2, d_{2}} x^{d_{2}}+c_{2, d_{2}-1} x^{d_{2}-1}+\cdots+c_{2,0}
\end{aligned}
$$

over $\mathbb{Z}$, with $f_{1} \neq \pm f_{2}$, both irreducible over $\mathbb{Z}$ and having content cont $f_{i}:=\operatorname{gcd}\left(c_{i, d_{i}}, \ldots, c_{i, 0}\right)$ equal to 1 ; we also choose an integer $m$ that is a common root modulo $n$ of $f_{1}$ and $f_{2}$. In our implementation this is the only step in which the SNFS and the GNFS differ: in the SNFS we use the special form of $n$ to pick these polynomials by hand. One polynomial will have very small coefficients compared to the coefficients of the polynomials we will use with the GNFS, where we search for a pair of polynomials with help of the computer. This makes SNFS faster than GNFS [Buhler et al. 1993, $\S 1]$. See Section 5 for a detailed description of the selection of the polynomials.

Let $\alpha_{i}$, for $i=1,2$, be a root of $f_{i}(x)$ in $\mathbb{C}$. Let $\mathbb{Q}_{n}$ denote the ring of rational numbers with denominator coprime to $n$. We want to find a nonempty set $\mathcal{S}$ of pairs $(a, b)$ of coprime integers such that both $\prod_{\mathscr{S}}\left(a-b \alpha_{1}\right)$ and $\prod_{\mathcal{S}}\left(a-b \alpha_{2}\right)$ are
squares - $\beta^{2}$ and $\gamma^{2}$, say - in $\mathbb{Q}_{n}\left[\alpha_{1}\right]$ and $\mathbb{Q}_{n}\left[\alpha_{2}\right]$, respectively. Applying to $\beta^{2}$ and $\gamma^{2}$ the two natural ring homomorphisms $\varphi_{i}: \mathbb{Q}_{n}\left[\alpha_{i}\right] \rightarrow \mathbb{Z} / n \mathbb{Z}$ determined by $\varphi_{i}\left(\alpha_{i}\right)=m \bmod n$ gives $\varphi_{1}\left(\beta^{2}\right) \equiv$ $\varphi_{2}\left(\gamma^{2}\right) \bmod n$. This yields $\varphi_{1}(\beta)^{2} \equiv \varphi_{2}(\gamma)^{2} \bmod n$. When $\varphi_{1}(\beta)$ and $\varphi_{2}(\gamma)$ are relatively prime to $n$, calculating $\operatorname{gcd}\left(n, \varphi_{1}(\beta)-\varphi_{2}(\gamma)\right)$ will yield a nontrivial factor of $n$ in at least half the cases.

For $\prod_{S}\left(a-b \alpha_{i}\right)$ to be a square in $\mathbb{Q}_{n}\left[\alpha_{i}\right]$, its norm $N\left(\prod_{\mathcal{S}}\left(a-b \alpha_{i}\right)\right)$ must be a square in $\mathbb{Q}$. Denote by $F_{i}(x, y)=y^{d_{i}} f_{i}(x / y) \in \mathbb{Z}[x, y]$ the homogeneous form of $f_{i}(x)$. From $N\left(a-b \alpha_{i}\right)=F_{i}(a, b) / c_{i, d_{i}}$ we can deduce that if the cardinality of the set $\mathcal{S}$ is even and if $\prod_{\mathcal{S}} F_{i}(a, b)$ is a square in $\mathbb{Z}$, then $N\left(\prod_{\mathscr{S}}\left(a-b \alpha_{i}\right)\right)$ is a square in $\mathbb{Q}$.

The algorithm searches for a pair $(a, b)$ of coprime integers such that both integers $F_{i}(a, b)$ factor completely over the prime numbers below some user-determined bounds $B_{i}$. We call such integers $F_{i}(a, b)$ smooth and such $(a, b)$-pairs relations. For a relation $\left(a_{j}, b_{j}\right)$ we can write

$$
\begin{align*}
& F_{1}\left(a_{j}, b_{j}\right)=\prod_{p \in \mathcal{K}_{1}} p^{e_{1}(j, p)} \\
& F_{2}\left(a_{j}, b_{j}\right)=\prod_{p \in \mathcal{K}_{2}} p^{e_{2}(j, p)} \tag{2.1}
\end{align*}
$$

where $e_{i}(j, p) \in \mathbb{N}$, for $i=1,2$, and where $\mathcal{K}_{1}$ and $\mathcal{K}_{2}$ contain -1 and the prime numbers below $B_{1}$ and $B_{2}$, respectively.

In order for $\prod_{\mathcal{S}} F_{i}(a, b)$ to be a square in $\mathbb{Z}$, every exponent $\sum_{\mathcal{S}} e_{i}(j, p)$ in

$$
\prod_{\left(a_{j}, b_{j}\right) \in \mathcal{S}} F_{i}\left(a_{j}, b_{j}\right)=\prod_{p \in \mathscr{K}_{i}} p^{\sum_{s} e_{i}(j, p)}
$$

should be even. Let $\boldsymbol{v}\left(a_{j}, b_{j}\right)$ be a vector of length $1+\left|\mathcal{K}_{1}\right|+\left|\mathcal{K}_{2}\right|$, constructed as follows: its first entry is 1 and the rest of $\boldsymbol{v}\left(a_{j}, b_{j}\right)$ is filled with all exponents $e_{1}(j, p)$ and $e_{2}(j, p)$ modulo 2 , in an order which is fixed for all $\left(a_{j}, b_{j}\right)$. If $\mathcal{S}$ is a subset of the relations such that $\sum_{(a, b) \in S} \boldsymbol{v}(a, b) \equiv \mathbf{0} \bmod 2$, then the cardinality of $\mathcal{S}$ is even and $\sum_{\mathcal{S}} e_{i}(j, p) \equiv$ $0 \bmod 2$ for $i=1,2$ and all $p \in \mathcal{K}_{i}$; hence both $N\left(\prod_{\mathcal{S}}\left(a-b \alpha_{i}\right)\right)$ are squares in $\mathbb{Q}$.

Unfortunately $N\left(\prod_{s}\left(a-b \alpha_{i}\right)\right)$ being a square in $\mathbb{Q}$ is not sufficient for $\prod_{\delta}\left(a-b \alpha_{i}\right)$ to be a square in $\mathbb{Q}_{n}\left[\alpha_{i}\right]$. By looking at what kind of $p$ divides $F_{i}(a, b)$, we will almost overcome this problem. For each prime number $p$ we define the set
$\mathcal{R}_{i}(p)=\left\{\left(r_{1}: r_{2}\right) \in \mathbf{P}^{1}\left(\mathbb{F}_{p}\right) \mid F_{i}\left(r_{1}, r_{2}\right) \equiv 0 \bmod p\right\}$,
where $\mathbf{P}^{1}\left(\mathbb{F}_{p}\right)$ denotes the projective line over $\mathbb{F}_{p}$. For $a$ and $b$ coprime, the integer $F_{i}(a, b)$ is divisible by a prime number $p$ if and only if $(a \bmod p: b \bmod$ $p) \in \mathcal{R}_{i}(p)$. Therefore the set $\mathcal{R}_{i}(p)$ is partitioning all $(a, b)$-pairs for which $p$ divides $F_{i}(a, b)$ according to $(a \bmod p: b \bmod p)$.

Next, for $C \in \mathbb{N}$, let $\mathcal{F}_{i}(C)$ be the set of pairs $\left(p,\left(r_{1}: r_{2}\right)\right)$, where $p$ is a prime less than $C$ and $\left(r_{1}: r_{2}\right) \in \mathcal{R}_{i}(p)$. Heuristically, $\left|\mathcal{F}_{i}(C)\right|$ is approximately the number of primes below $C$ [Lang 1970, Chapter VIII, $\S 4] . \mathcal{F}_{1}\left(B_{1}\right)$ and $\mathcal{F}_{2}\left(B_{2}\right)$ are called the factor bases. We now can write (2.1) as

$$
\begin{equation*}
F_{i}\left(a_{j}, b_{j}\right)= \pm \prod_{\left(p,\left(r_{1}: r_{2}\right)\right) \in \mathcal{F}_{i}\left(B_{i}\right)} p^{e_{i}\left(j, p, r_{1}, r_{2}\right)} \tag{2.3}
\end{equation*}
$$

for $i=1,2$, where $e_{i}\left(j, p, r_{1}, r_{2}\right)=e_{i}(j, p)$ if

$$
\left(a_{j} \bmod p: b_{j} \bmod p\right)=\left(r_{1}: r_{2}\right)
$$

and 0 otherwise.
In order for $\prod_{s}\left(a-b \alpha_{i}\right)$ to be a square in $\mathbb{Q}_{n}\left[\alpha_{i}\right]$, every exponent $\sum_{\delta} e_{i}\left(j, p, r_{1}, r_{2}\right)$ in

$$
\prod_{\left(a_{j}, b_{j}\right) \in \mathcal{S}} F_{i}\left(a_{j}, b_{j}\right)= \pm \prod_{\left(p,\left(r_{1}: r_{2}\right)\right) \in \mathcal{F}_{i}\left(B_{i}\right)} p^{\sum_{\mathcal{E}} e_{i}\left(j, p, r_{1}, r_{2}\right)}
$$

should be even. Let $\boldsymbol{v}\left(a_{j}, b_{j}\right)$ be a vector of length $1+\left|\mathcal{F}_{1}\left(B_{1}\right)\right|+\left|\mathcal{F}_{2}\left(B_{2}\right)\right|$ containing 1 and the values of $e_{1}\left(j, p, r_{1}, r_{2}\right) \bmod 2$ and $e_{2}\left(j, p, r_{1}, r_{2}\right) \bmod 2$, in an order that is fixed for all relations $\left(a_{j}, b_{j}\right)$. A nonempty subset $\mathcal{S}$ of relations such that

$$
\sum_{\delta} \boldsymbol{v}(a, b) \equiv \mathbf{0} \bmod 2
$$

is almost sufficient to ensure that $\prod_{\delta}\left(a-b \alpha_{i}\right)$ be a square in $\mathbb{Q}_{n}\left[\alpha_{i}\right]$ for $i=1,2$ [Buhler et al. 1993, $\S 12.7]$. That it is not totally sufficient is only partly caused by the fact that we only forced the
product $\prod_{S}\left|F_{i}(a, b)\right|$ to be a square in $\mathbb{Z}$. We can see that it is not totally sufficient from the following example: In the field $\mathbb{Q}(\sqrt{3})$ generated by a root of the polynomial $f(x)=x^{2}-3$, the element $2+\sqrt{3}$ has norm $F(2,-1)=1$. So all exponents $e_{1}\left(j, p, r_{1}, r_{2}\right)$ and $e_{2}\left(j, p, r_{1}, r_{2}\right)$ will be zero. Furthermore $\boldsymbol{v}(2,-1)+\boldsymbol{v}(1,0) \equiv \mathbf{0} \bmod 2$. But the square root of $2+\sqrt{3}$ is $(\sqrt{6}+\sqrt{2}) / 2$, which is not an element of $\mathbb{Q}(\sqrt{3})$.

The small gap between being almost a square and being practically certainly a square is overcome by using quadratic characters, following an idea of Adleman [1991]. For $\mathcal{S}$ a set consisting of pairs $(a, b)$ of coprime integers, let $\prod_{\delta}\left(a-b \alpha_{i}\right)$ be a square in $\mathbb{Q}_{n}\left[\alpha_{i}\right]$, and let $q$ be an odd prime number not dividing $c_{1, d_{1}} c_{2, d_{2}}$. If $\left(s_{1}: s_{2}\right) \in \mathcal{R}_{i}(q)$ is such that $f_{i}^{\prime}\left(s_{1} s_{2}^{-1} \bmod q\right) \not \equiv 0 \bmod q$ and $(a \bmod q:$ $b \bmod q) \neq\left(s_{1}: s_{2}\right)$ for all $(a, b) \in \mathcal{S}$, then

$$
\begin{equation*}
\prod_{(a, b) \in s}\left(\frac{a-b\left(s_{1} s_{2}^{-1} \bmod q\right)}{q}\right)=1 \tag{2.4}
\end{equation*}
$$

where $\left(\frac{v}{w}\right)$ denotes the Legendre symbol [Buhler et al. 1993, § 8, § 12.7]. We use this by taking for each polynomial several primes $q$ larger than $B_{i}$ and not dividing $c_{i, d_{i}}$, together with an element $\left(s_{1}: s_{2}\right) \in$ $\mathcal{R}_{i}(q)$ such that $f_{i}^{\prime}\left(s_{1} s_{2}^{-1} \bmod q\right) \not \equiv 0 \bmod q$. Since $q>B_{i}$ we have $(a \bmod q: b \bmod q) \neq\left(s_{1}: s_{2}\right)$ for all relations $(a, b)$. Append to the vector $\boldsymbol{v}(a, b)$ for all pairs $\left(q,\left(s_{1}: s_{2}\right)\right)$ a 0 if

$$
\left(\frac{a-b\left(s_{1} s_{2}^{-1} \bmod q\right)}{q}\right)=1
$$

and a 1 otherwise. Now a nonempty subset $\mathcal{S}$ of all relations such that $\sum_{\delta} \boldsymbol{v}(a, b) \equiv \mathbf{0} \bmod 2$ guarantees that (2.4) holds for all chosen primes $q$ together with their elements $\left(s_{1}: s_{2}\right) \in \mathcal{R}_{i}(q)$. Taking enough quadratic characters - we took 32 per polynomial - makes it practically certain that both $\prod_{s}\left(a-b \alpha_{i}\right)$ are squares in $\mathbb{Q}_{n}\left[\alpha_{i}\right]$. The counterexample given earlier could have been caught with the use of quadratic characters: take $q=11$ and $\left(s_{1}: s_{2}\right)=(5,1) \in \mathcal{R}(11)$. The Legendre symbol becomes $\left(\frac{2+5}{11}\right)$, which is -1 .

If $Q$ is the total number of quadratic characters used, then a nonempty subset $\mathcal{S}$ such that $\sum_{\delta} \boldsymbol{v}(a, b) \equiv \mathbf{0} \bmod 2$ can always be found if the number of relations exceeds

$$
1+\left|\mathfrak{F}_{1}\left(B_{1}\right)\right|+\left|\mathcal{F}_{2}\left(B_{2}\right)\right|+Q .
$$

## 3. OUTLINE OF THE IMPLEMENTATION

The implementation can be divided into five stages. In the first stage we select the polynomials $f_{1}(x)$ and $f_{2}(x)$ in $\mathbb{Z}[x]$, and the integer $m$ such that $m$ is a common root of $f_{1}(x)$ and $f_{2}(x)$ modulo $n$. We also choose the sieving region-that is, the collection of $(a, b)$-pairs for which both $F_{i}(a, b)$ are checked for smoothness - and, for each polynomial, a factor base bound $B_{i}$.

The second stage, the sieving in which the relations are found, is the most time-consuming. In this implementation a relation is a pair $(a, b)$ from the sieving region such that both $F_{i}(a, b)$ factor completely over the primes below $B_{i}$, except for at most two large prime numbers, which should be between $B_{i}$ and a large prime bound $L_{i}$. By using lattice sieving [Pollard 1993b] - a special form of which will be desribed in Section 6 - one of the two integers $F_{i}(a, b)$ is allowed to have three primes between $B_{i}$ and $L_{i}$. The product in (2.3) is taken over $\mathcal{F}\left(L_{i}\right)$, and the vectors $\boldsymbol{v}(a, b)$ have to be adapted accordingly.

This is followed by a filtering stage with the purpose of reducing the amount of data. Here some relations are eliminated and others are grouped into relation-sets.

In the fourth stage, we construct a matrix by taking the vectors $\boldsymbol{v}(a, b)$ and the vectors

$$
\sum_{(a, b) \in \mathcal{V}} \boldsymbol{v}(a, b)
$$

for all remaining relations and relation-sets $\mathcal{V}$ as columns. Finding a nonempty set $\mathcal{S}$ such that

$$
\sum_{\delta} \boldsymbol{v}(a, b) \equiv \mathbf{0} \bmod 2
$$

is the same as calculating a nontrivial vector from the null space of this matrix over $\mathbb{F}_{2}$. For huge sparse matrices the best known methods are iterative ones, such as the block Lanczos algorithm [Montgomery 1995]. The output of this stage is a subset $\mathcal{S}$ of the relations such that both $\prod_{\mathcal{S}}\left(a-b \alpha_{1}\right)$ and $\prod_{\S}\left(a-b \alpha_{2}\right)$ are squares $\beta^{2}$ and $\gamma^{2}$ in $\mathbb{Q}_{n}\left[\alpha_{1}\right]$ and $\mathbb{Q}_{n}\left[\alpha_{2}\right]$, respectively.

The final stage consists of extracting the square roots $\beta$ and $\gamma$. This is done by a new algorithm, developed by Montgomery [1994] and also iterative. Successive approximations are found, leaving over "smaller" remainders of which we have to extract the square root. If the remainder is small enough we use a conventional method. Finally we apply the homomorphisms $\varphi_{1}$ and $\varphi_{2}$ to the square roots $\beta$ and $\gamma$, respectively, and calculate the gcd of $n$ and $\varphi_{1}(\beta)-\varphi_{2}(\gamma)$, which will split $n$ into two nontrivial factors in at least half of the cases.

## 4. FREE RELATIONS

Denote the order of the Galois group of $f_{1}(x) f_{2}(x)$ by $g$. For approximately $1 / g$ of the primes $q<$ $\min \left(L_{1}, L_{2}\right)$, both polynomials $F_{i}(x, y)$ split into $d_{i}$ linear factors modulo $q$ [Frobenius 1896, $\S 2$, Theorem 1; Neukirch 1992, p. 566]:

$$
\begin{equation*}
F_{i}(x, y)=c_{i, d_{i}} \prod_{j=1}^{d_{i}}\left(r_{2}^{(j)} x-r_{1}^{(j)} y\right) \bmod q \tag{4.1}
\end{equation*}
$$

If such a prime $q$ does not divide the discriminants of $f_{1}(x)$ or $f_{2}(x)$ (and therefore both polynomials $F_{i}(x, y)$ split into $d_{i}$ different linear factors modulo $q$ ) and if $q$ does not divide $c_{1, d_{1}} \cdot c_{2, d_{2}}$, we call $q$ a free prime. This terminology comes from the fact that we can select such primes that are smaller than $\min \left(B_{1}, B_{2}\right)$ without extra effort when calculating the factor bases $\mathcal{F}_{i}\left(B_{i}\right)$. They are said to give rise to free relations because we now require $\left(\prod_{p \in \mathcal{T}} p\right)\left(\prod_{(a, b) \in \mathcal{S}}\left(a-b \alpha_{i}\right)\right)$ to be a square in $\mathbb{Q}_{n}\left[\alpha_{i}\right]$, for $i=1,2$, where $\mathfrak{T}$ is a suitably chosen subset of the set of free primes. With $N(p)=p^{d_{i}}$, we have

$$
\begin{aligned}
& N\left(\left(\prod_{p \in \mathcal{T}} p\right)\left(\prod_{(a, b) \in \mathcal{S}}\left(a-b \alpha_{i}\right)\right)\right) \\
&=\left(\prod_{p \in \mathcal{T}} p^{d_{i}}\right)\left(\prod_{(a, b) \in \mathcal{S}} \frac{F_{i}(a, b)}{c_{i, d_{i}}}\right)
\end{aligned}
$$

which represents a square in $\mathbb{Q}$ if $|\mathcal{S}|$ is even and $\left(\prod_{p \in \mathcal{T}} p^{d_{i}}\right)\left(\prod_{(a, b) \in \mathcal{S}} F_{i}(a, b)\right)$ is a square in $\mathbb{Z}$. As we partitioned the primes $p$ dividing $\prod_{(a, b) \in S} F_{i}(a, b)$ according to the roots $(a \bmod p: b \bmod p) \in \mathcal{R}_{i}(p)$, we consider $p^{d_{i}}$ as the product of one factor $p$ for every root $\left(r_{1}: r_{2}\right) \in \mathcal{R}_{i}(p)$. We associate with every free prime $p<\min \left(L_{1}, L_{2}\right)$ a vector $\boldsymbol{v}(p)$ of length $1+\left|\mathcal{F}_{1}\left(L_{1}\right)\right|+\left|\mathcal{F}_{2}\left(L_{2}\right)\right|+Q$, which contains a one for every $\left(p,\left(r_{1}^{(j)}: r_{2}^{(j)}\right)\right)$ occurring in (4.1) for both polynomials, and for every quadratic character $\left(q,\left(s_{1}: s_{2}\right)\right)$ for which $\left(\frac{p}{q}\right)=-1$. The rest is filled with zeros. We will look for $\mathcal{T}$ and $\mathcal{S}$ such that $\sum_{p \in \mathcal{T}} \boldsymbol{v}(p)+\sum_{(a, b) \in \mathcal{S}} \boldsymbol{v}(a, b) \equiv \mathbf{0} \bmod 2$.

## 5. CHOICE OF THE POLYNOMIALS

The conjectured running time for the application of the SNFS to a number of the form $n=c_{1} r^{t}+c_{2} s^{u}$ depends on the size of $n$. If only small factors of $n$ are known, the SNFS algorithm is certainly the best one to use. If already a substantial nonalgebraic factor of $n$ is known, the GNFS or the MPQS might be faster.

Using the SNFS for a factor $n$ of an integer $c_{1} r^{t}+c_{2} s^{u}$ with $\operatorname{gcd}\left(c_{1} r, c_{2} s\right)=1$, we pick the two possibly nonmonic polynomials by hand. Select a small positive integer $d_{1}$ - usually 4 or 5 - which will be the degree of $f_{1}(x)$. Write $t=d_{1} t^{\prime}+t^{\prime \prime}$ and $u=d_{1} u^{\prime}+u^{\prime \prime}$ with $t^{\prime \prime}, u^{\prime \prime} \in\left\{0,1, \ldots, d_{1}-1\right\}$. In practice $f_{1}(x):=c_{2} s^{u^{\prime \prime}} x^{d_{1}}+c_{1} r^{t^{\prime \prime}}$ is irreducible over $\mathbb{Z}$, and $f_{1}(x), f_{2}(x):=r^{t^{\prime}} x-s^{u^{\prime}}$, and $m:=$ $s^{u^{\prime}} r^{-t^{\prime}} \bmod n$ satisfy the requirements mentioned in Section 2. If $f_{1}(x)$ is not irreducible, a nontrivial factor of $f_{1}$ is likely to give rise to a nontrivial factor of $n$, and otherwise $f_{1}$ can be replaced by a suitable factor. (This is also applicable in the case of the GNFS.) An algorithm to test whether a polynomial is irreducible and to factor it if it is not can be found in [Lenstra et al. 1982]. If we
encounter a polynomial $f_{i}(x)$ with cont $f_{i}(x) \neq 1$, we can divide all coefficients of $f_{i}(x)$ by the content, assuming that cont $f_{i}(x)$ and $n$ are relatively prime. Using the SNFS we sometimes find better pairs of polynomials, together with a value for $m$, by trying to factor a multiple of $n$. Examples can be found in the last section of this article.

Using the GNFS one can find two polynomials by the base $m$ method. Select a small positive integer $d_{1}$ - usually 4 or 5 - which will again be the degree of $f_{1}(x)$. Set $m=\left\lfloor n^{1 / d_{1}}\right\rfloor$ and write $n$ in base $m$ as

$$
n=c_{d_{1}} m^{d_{1}}+c_{d_{1}-1} m^{d_{1}-1}+\cdots+c_{0}
$$

with $0 \leq c_{i}<m$. Now

$$
f_{1}(x)=c_{d_{1}} x^{d_{1}}+c_{d_{1}-1} x^{d_{1}-1}+\cdots+c_{0}
$$

and $f_{2}(x)=x-m$ satisfy the requirements. This method implies $c_{d_{1}}=1$ [Buhler et al. 1993, § 3]. In [Buhler et al. 1993, § 12.2] one can find slightly better variants of this method, resulting in a linear and a higher-degree polynomial with leading coefficients possibly larger than one and possibly negative coefficients. For these variants the polynomial coefficients are $\mathcal{O}\left(n^{1 /\left(d_{1}+1\right)}\right)$.

The task is to find suitable polynomials $f_{1}$ and $f_{2}$, factor base bounds $B_{1}, B_{2}$, large prime bounds $L_{1}, L_{2}$, and a sieving region. For a good choice four characteristics of the polynomials should be taken into account. First, the maximal values of $\left|F_{i}(a, b)\right|$ should be small, making them more likely to be smooth over the primes below $B_{i}$. Secondly, when a polynomial has many real roots, more ra$\operatorname{tios} a / b$ will be near a root and more values $F_{i}(a, b)$ are expected to be small. As a refinement of this characteristic we can look at the absolute value of the real roots. A polynomial having a real root near $\max |a| / \max |b|$ is a good choice. The importance of this characteristic is made clear in Figure 1. Thirdly, polynomials that have many roots modulo (preferably different) small primes are preferred over ones that do not. This enlarges the probability that $F_{i}(a, b)$ is small after dividing it by
these small prime numbers, making it more likely to be smooth over the primes below $B_{i}$. Finally, it is better to choose polynomials for which the order of the Galois group of $f_{1}(x) f_{2}(x)$ is small, since we saw in the previous section that they provide more free relations. With these criteria in mind we select the pair of polynomials which is expected to be the best.


FIGURE 1. Number of relations found for 60000 $a$-values and $8625 b$-values for the factorization of a 119-digit factor of $3^{319}-1$ (see Section 10 for details). One polynomial is $f_{1}(x)=x^{5}+x^{4}-4 x^{3}-$ $3 x^{2}+3 x+1$, with 5 real roots. The five ridges indicate a higher yield for pairs $(a, b)$ with $a / b$ near a root.

We experimented with a choice of two quadratic polynomials selected according to ideas of Montgomery [Buhler et al. 1994]. He observed that $f_{1}(x)=c_{1,2} x^{2}+c_{1,1} x+c_{1,0}$ and $f_{2}(x)=c_{2,2} x^{2}+$ $c_{2,1} x+c_{2,0} \in \mathbb{Z}[x]$ have a common root $m$ modulo $n$ if and only if the vectors $\boldsymbol{a}=\left(c_{1,0}, c_{1,1}, c_{1,2}\right)^{T}$ and $\boldsymbol{b}=\left(c_{2,0}, c_{2,1}, c_{2,2}\right)^{T}$ are orthogonal to $\left(1, m, m^{2}\right)^{T}$ over $\mathbb{Z} / n \mathbb{Z}$ using the standard inner product. Suppose $f_{1}(x)$ and $f_{2}(x)$ are irreducible over $\mathbb{Z}$, have content 1 , and do not satisfy $f_{1}(x)= \pm f_{2}(x)$. As will be explained further on, we can find in practice $\boldsymbol{a}$ and $\boldsymbol{b}$ of which the coefficients are appoximately $\mathcal{O}\left(n^{1 / 4}\right)$, so the space orthogonal to $\boldsymbol{a}$ and $\boldsymbol{b}$ has rank 1 (both over $\mathbb{Z}$ and over $\mathbb{Z} / n \mathbb{Z}$ ). If
$\boldsymbol{c}=\boldsymbol{a} \times \boldsymbol{b}$ (cross product), then $\boldsymbol{c}$ must be a multiple of $\left(1, m, m^{2}\right)^{T}$ over $\mathbb{Z} / n \mathbb{Z}$. The fact that $f_{1}(x)$ and $f_{2}(x)$ are not multiples of each other ensures that $\boldsymbol{c}$ is not the zero vector. If $\boldsymbol{c}=\left(c_{0}, c_{1}, c_{2}\right)^{T}$, then $c_{0}, c_{1}, c_{2}$ is a geometric progression in $\mathbb{Z} / n \mathbb{Z}$. It is not a geometric progression over $\mathbb{Z}$, since then $f_{1}(x)$ and $f_{2}(x)$ would have a common factor $x-m$ over $\mathbb{Z}$.

Montgomery's algorithm for finding $f_{1}(x)$ and $f_{2}(x)$ reverses this construction and starts with a vector $\boldsymbol{c}=\left(c_{0}, c_{1}, c_{2}\right)^{T} \in \mathbb{Z}^{3}$, where $c_{0}, c_{1}, c_{2}$ is a geometric progression with ratio $m$ over $\mathbb{Z} / n \mathbb{Z}$, but not over $\mathbb{Z}$. The vector $\boldsymbol{c}$ can be constructed as follows: for $p$ prime such that $p<\sqrt{n}$ and $n$ a quadratic residue modulo $p$, choose $c_{1}$ such that $c_{1}^{2} \equiv n \bmod p$ and $\left|c_{1}-n^{1 / 2}\right| \leq p / 2$. The elements of $\boldsymbol{c}=\left(p, c_{1},\left(c_{1}^{2}-n\right) / p\right)^{T}$ form a geometric progression with ratio $c_{1} / p$ over $\mathbb{Z} / n \mathbb{Z}$, not over $\mathbb{Z}$. Furthermore $c_{i}=\mathcal{O}\left(n^{1 / 2}\right)(i=0,1,2)$. Take $s \in \mathbb{Z} / p \mathbb{Z}$ such that $c_{1} s \equiv 1 \bmod p$. With $c_{2}=\left(c_{1}^{2}-n\right) / p$, the vectors

$$
\boldsymbol{a}^{\prime}=\left(\begin{array}{c}
c_{1} \\
-p \\
0
\end{array}\right), \quad \boldsymbol{b}^{\prime}=\left(\begin{array}{c}
\left(c_{1}\left(c_{2} s \bmod p\right)-c_{2}\right) / p \\
-\left(c_{2} s \bmod p\right) \\
1
\end{array}\right)
$$

are both orthogonal to $\boldsymbol{c}$. From $\boldsymbol{a}^{\prime} \times \boldsymbol{b}^{\prime}=-\boldsymbol{c}$ and $\operatorname{gcd}\left(c_{0}, c_{1}, c_{2}\right)=1$ we deduce that $\boldsymbol{a}^{\prime}$ and $\boldsymbol{b}^{\prime}$ span the sublattice of $\mathbb{Z}^{3}$ orthogonal to $\boldsymbol{c}$. Denote by $(\boldsymbol{a}, \boldsymbol{b})$ the inner product of $\boldsymbol{a}$ and $\boldsymbol{b}$, and remember that $\boldsymbol{a}(\boldsymbol{a}, \boldsymbol{b}) /(\boldsymbol{a}, \boldsymbol{a})$ is the projection of $\boldsymbol{b}$ on $\boldsymbol{a}$. By reducing the basis $\left\{\boldsymbol{a}^{\prime}, \boldsymbol{b}^{\prime}\right\}$, one can find "small" vectors $\boldsymbol{a}$ and $\boldsymbol{b}$ with $\{\boldsymbol{a}, \boldsymbol{b}\}$ a basis of the sublattice of $\mathbb{Z}^{3}$ orthogonal to $\boldsymbol{c}$, such that

$$
\left|\frac{(\boldsymbol{a}, \boldsymbol{b})}{(\boldsymbol{a}, \boldsymbol{a})}\right| \leq \frac{1}{2} \quad \text { and } \quad\left|\frac{(\boldsymbol{a}, \boldsymbol{b})}{(\boldsymbol{b}, \boldsymbol{b})}\right| \leq \frac{1}{2} .
$$

The angle $\theta$ between these vectors will be between $60^{\circ}$ and $120^{\circ}$. Since the surface of the parallelogram spanned by $\boldsymbol{a}$ and $\boldsymbol{b}$ is both equal to $\|\boldsymbol{a} \times \boldsymbol{b}\|$ and $\|\boldsymbol{a}\| \cdot\|\boldsymbol{b}\| \sin \theta$, we have

$$
\begin{equation*}
\|\boldsymbol{a}\| \cdot\|\boldsymbol{b}\|=\frac{\|\boldsymbol{c}\|}{\sin \theta} \leq \frac{2\|\boldsymbol{c}\|}{\sqrt{3}}=\mathcal{O}(\|\boldsymbol{c}\|)=\mathcal{O}\left(n^{1 / 2}\right) . \tag{5.1}
\end{equation*}
$$

In practice both $\|\boldsymbol{a}\|$ and $\|\boldsymbol{b}\|$ are $\mathcal{O}\left(n^{1 / 4}\right)$. For different values of $p$ we will get a different pair of polynomials. The program findquad tries to find two polynomials each having two real roots, many roots modulo small primes, and such that the integral of $\sum_{i=1,2} \log \left|F_{i}(x, y)\right|$ is small, where $(x, y)$ runs through the sieving region for $(a, b)$.
When using line sieving (a special form of lattice sieving, explained in Section 6), we like to use a large range of $a$-values, say $|a|<M$ and only $b=1$. To try to induce $F_{i}(a, 1)=c_{i, 2} a^{2}+$ $c_{i, 1} a+c_{i, 0}$ to be smooth over the prime numbers below $B_{i}$, we would prefer $c_{i, 2}=\mathcal{O}\left(n^{1 / 4} / M\right), c_{i, 1}=$ $\mathcal{O}\left(n^{1 / 4}\right)$, and $c_{i, 0}=\mathcal{O}\left(n^{1 / 4} M\right)$ rather than all of them being $\mathcal{O}\left(n^{1 / 4}\right)$. We achieve this by first choosing $c_{0}=p=\mathcal{O}(\sqrt{n} / M)$, whence $c_{1}=\mathcal{O}(\sqrt{n})$ and $c_{2}=\mathcal{O}(\sqrt{n} M)$. The resulting coefficients of $\boldsymbol{a}^{\prime}=\left(a_{0}^{\prime}, a_{1}^{\prime}, a_{2}^{\prime}\right)$ and $\boldsymbol{b}^{\prime}=\left(b_{0}^{\prime}, b_{1}^{\prime}, b_{2}^{\prime}\right)$ have approximately the right ratio. To keep this ratio while reducing the basis, we reduce the vectors $\boldsymbol{a}^{\prime \prime}=$ $\left(a_{0}^{\prime}, a_{1}^{\prime} M, a_{2}^{\prime} M^{2}\right)$ and $\boldsymbol{b}^{\prime \prime}=\left(b_{0}^{\prime}, b_{1}^{\prime} M, b_{2}^{\prime} M^{2}\right)$ instead. Note that $\boldsymbol{a}^{\prime \prime} \times \boldsymbol{b}^{\prime \prime}=M \boldsymbol{c}^{\prime \prime}$ with $\boldsymbol{c}^{\prime \prime}=\left(c_{0} M^{2}, c_{1} M, c_{2}\right)$, which is still a geometric progression with ratio $c_{1} / p M$ over $\mathbb{Z} / n \mathbb{Z}$, not over $\mathbb{Z}$. Using (5.1) with $\boldsymbol{c}=M \boldsymbol{c}^{\prime \prime}$ we find that the resulting vectors $\boldsymbol{a}=$ $\left(a_{0}, a_{1} M, a_{2} M^{2}\right)$ and $\boldsymbol{b}=\left(b_{0}, b_{1} M, b_{2} M^{2}\right)$ will be both $\mathcal{O}\left(n^{1 / 4} M\right)$. Using $f_{1}(x)=a_{2} x^{2}+a_{1} x+a_{0}$ and $f_{2}(x)=b_{2} x^{2}+b_{1} x+b_{0}$ results in the desired orders of the coefficients of $f_{1}$ and $f_{2}$.

We also have to choose the factor base bounds $B_{i}$, the large prime bounds $L_{i}$, and the sieving region. In the experiments described in Section 10, where we factored numbers in the 98-162 digits range with the SNFS and numbers in the $87-$ 107 digits range with the GNFS, we used factor base bounds between $5 \cdot 10^{5}$ and $2.9 \cdot 10^{6}$ and large prime bounds between $12 \cdot 10^{6}$ and $4 \cdot 10^{7}$. When using classical sieving, the sieving region was a rectangle for which we took $a$ in a subinterval of $\left[-2 \cdot 10^{6}, 2 \cdot 10^{6}\right]$ and $b$ between 1 and some upper bound, in our experiments between $16 \cdot 10^{3}$ and $48 \cdot 10^{4}$. First we chose the factor base bounds and generated the corresponding factor bases $\mathcal{F}_{i}\left(B_{i}\right)$, as described in the next section. Then we chose
the large prime bounds $L_{i}$ and fixed a range of $a$ values. For all these $a$-values and a few $b$-values, preferably equidistributed over the expected range of $b$-values, we checked whether the $(a, b)$-pair is a relation, in a way we will describe in the next section. Allowing $F_{i}(a, b)$ to contain two large primes between $B_{i}$ and $L_{i}$, instead of demanding it to be smooth over the primes below $B_{i}$, we increase the probability that $(a, b)$ is a relation. On the other hand, $F_{i}(a, b)$ is now factored over $\mathcal{F}_{i}\left(L_{i}\right)$ instead of $\mathcal{F}_{i}\left(B_{i}\right)$, which enlarges the number of relations needed too. In practice this adjustment has shown to be useful. In our experiments we needed approximately $0.8 \cdot\left\{\pi\left(L_{1}\right)+\pi\left(L_{2}\right)-\pi\left(\min \left(L_{1}, L_{2}\right)\right) / g\right\}$ relations.

From the number of relations we got for these few $b$-values we could estimate the range of $b$-values needed and from the time the experiment took we could estimate the time needed for the whole sieving step. In this way we selected a good combination of pair of polynomials, factor base bounds, large prime bounds and sieving region.

## 6. THE SIEVING

The sieving is the part of the algorithm during which we collect the relations. Before we start the sieving we have to generate the factor bases $\mathcal{F}_{i}\left(B_{i}\right)$ defined on page 233, just before (2.3). If we identify $\mathbf{P}^{1}\left(\mathbb{F}_{p}\right)$ with $\mathbb{F}_{p} \cup\{\infty\}$ by identifying $\left(r_{1}: r_{2}\right)$ with $r_{1} / r_{2}$, then $\mathcal{R}_{i}(p)$ (defined in (2.2)) consists of those $r=r_{1} / r_{2} \in \mathbb{F}_{p}$ for which $f_{i}(r) \equiv$ $0 \bmod p$, together with $\infty$ if $c_{i, d_{i}} \equiv 0 \bmod p$. The program rootfinder finds for both polynomials $f_{i}(x)$ and for all primes $p$ below $B_{i}$ all roots modulo $p$. Repeated roots appear only once in the list. When the prime $p$ divides the leading coefficient $c_{i, d_{i}}$, rootfinder includes the projective root $(1,0)$, which it represents by $p$. We recall that for $a$ and $b$ coprime, $p \mid F_{i}(a, b)$ if and only if $(a \bmod p$ : $b \bmod p) \in \mathcal{R}_{i}(p)$. In terms of the roots of $f_{i}(x)$ this means that for $a$ and $b$ coprime, $p \mid F_{i}(a, b)$ if and only if $\left(a \equiv b r \bmod p\right.$ and $\left.f_{i}(r) \equiv 0 \bmod p\right)$ or ( $p \mid c_{i, d_{i}}$ and $p \mid b$ ).

Two ways of sieving have been implemented: the "classical" sieve [Lenstra et al. 1993b, §4; Buhler et al. 1993, §12], and line sieving, a special form of lattice sieving [Pollard 1993b].

In the classical way of sieving we first choose the $a$-interval and the $b$-interval. We start sieving with $b=1$ and augment $b$ until we reach its upper bound. The program gnfs estimates the maximum value of $F_{i}(a, b)$ over all values of $a$ and $b$ for both polynomials. The polynomial for which this estimate is larger is sieved first. Probably fewer pairs $(a, b)$ will have a smooth value of $F_{i}(a, b)$ for this polynomial, so fewer pairs have to be stored. Furthermore this largest value is used to decide upon the base of the logarithm, which we choose in such a way that the $\log$ of the maximum fits in one byte. Suppose we start sieving with polynomial $f_{j}$.

To sieve for the first polynomial $f_{j}$ we fix $b$ and initialize to zero an array that contains one byte per $a$-value. For every prime $p<B_{j}$ and every $r$ with $f_{j}(r) \equiv 0 \bmod p$ we add $[\log p]$ (where [ • ] is the nearest integer function) to all array elements corresponding with $a \equiv b r \bmod p$. For every prime $p<B_{j}$ with $p \mid c_{j, d_{j}}$ and $p \mid b$ we add $[\log p]$ to every array element. Then we split the $a$-interval recursively in subintervals until the value of $c_{j}(a, b)=F_{j}(a, b) / L_{j}^{2}$ does not vary more than a prescribed amount within a subinterval. If the value of an array element is close enough to $\log c_{j}(a, b)$, then $F_{j}(a, b)$ is potentially smooth and we store the value of $a$. Now the same sieving process takes place for the other polynomial $f_{3-j}$. If for a pair $(a, b)$ both $F_{1}(a, b)$ and $F_{2}(a, b)$ are potentially smooth - $(a, b)$ is now called a candidate relation -, we use trial division (where we first test if $a \equiv b r \bmod p$ before applying an expensive multiple precision division of $F_{i}(a, b)$ on $p$ ) to extract all factors below $B_{i}$ from $F_{i}(a, b)$, for $i=1,2$. This is necessary, since during the sieving we use rounded logarithms and other techniques, which not only make the sieving faster, but also make the final value in the array elements less accurate.
(In [Golliver et al. 1994] experiments were made with repeating the sieving procedure once again,
instead of using trial division. The candidate relations are marked in the sieving array. In a second sieving round the primes $p$ themselves are stored instead of adding $[\log p]$ to the array elements for the candidate relations. Next the integers $F_{1}(a, b)$ and $F_{2}(a, b)$ are calculated for the candidate relations and the stored primes are divided out. This approach costs more memory, but is likely faster.)

By increasing $p$, and comparing the sieved logarithms with the sum of the logarithms of primes divided out of $F_{i}(a, b)$ during trial division so far, one can sometimes skip an interval of primes. If after the trial division there remains a composite part smaller than $L_{i}^{2}$, we try to factor it first using SQUFOF, and if that fails using Pollard Rho [Riesel 1985, pp. 191-198, 174-183]. A pair ( $a, b$ ) is a relation if both $F_{i}(a, b)$ factor over the primes below $B_{i}$ except for at most two large primes between $B_{i}$ and $L_{i}$. It is stored together with the primes dividing $F_{i}(a, b)$ that exceed some user-determined printing bounds $W_{i}$, where $i=1,2$. With these bounds $W_{i}$ one can monitor the amount of output of the gnfs program. They should be chosen in such a way that it fits in the available disk space.

Using the lattice sieve, we only sieve over pairs $(a, b)$ of which we know that one $F_{i}(a, b)$, say for $i=j \in\{1,2\}$, is divisible by a special large prime between $L^{(l)}$ and $L^{(u)}$, which are the user-chosen lower and upper bound for the large primes, respectively. The advantage is that the remaining part of $F_{j}(a, b)$ is more likely to be smooth. On the other hand we will miss the relations for which both $F_{i}(a, b)$ are smooth over the primes below $L^{(l)}$. For the implementation of the lattice sieve we use an extra feature implemented in the classical way of sieving. There we have a possibility of sieving over a sublattice of the ( $a, b$ )-pairs. We can choose an integral, nonsingular matrix $\boldsymbol{M}$ and sieve over pairs $(a, b)$ of the form:

$$
\binom{a}{b}=M\binom{x}{y},
$$

while the program sieves over $x$ and $y$. This is done by substituting the expressions of $a$ and $b$ in
terms of $x$ and $y$ in both $F_{i}(a, b)$ resulting in new polynomials $G_{i}(x, y)$, which are now the polynomials whose values should be smooth. Of course the roots of the polynomials $F_{i}$ have to be adapted to the roots of the polynomials $G_{i}$. When a pair $(x, y)$ is a relation, the corresponding pair $(a, b)$, together with the primes dividing $G_{i}(x, y)$ and exceeding $W_{i}$, are stored.

The lattice sieve sieves for every prime $q$ in the range $\left[L^{(l)}, L^{(u)}\right]$, for a fixed value of $b$ over all roots $\left(r_{1}: r_{2}\right) \in\left\{\mathcal{R}_{1}(q) \cup \mathcal{R}_{2}(q)\right\}$ with $r_{2} \neq 0$. When sieving over a root $\left(r_{1}: r_{2}\right)$ of $\mathcal{R}_{j}(q)$ we sieve only over the $a$-values with $a \equiv b r_{1} r_{2}^{-1} \bmod q$, thus guaranteeing that $F_{j}(a, b)$ is divisible by $q$. This is the same as using a matrix

$$
\boldsymbol{M}=\left(\begin{array}{cc}
q & r_{1} r_{2}^{-1} \bmod q \\
0 & 1
\end{array}\right),
$$

with $y$ fixed to $b$ and $x$ in an interval such that $q x+$ $b\left(r_{1} r_{2}^{-1} \bmod q\right)$ just fits in the $a$-interval. (Note that, when we compare our notation with that used in [Pollard 1993b], we have $V_{1}=(q, 0)$ and $V_{2}=\left(r_{1} r_{2}^{-1} \bmod q, 1\right)$, and that we are applying the "sieving by rows" strategy.) $G_{j}(x, y) / q$ should be smooth over the primes below $B_{j}$, except for at most two large primes between $B_{j}$ and $q$. The other $G_{3-j}(x, y)$ should be smooth over the primes below $B_{3-j}$, except for at most two large primes between $B_{3-j}$ and $q$. Not allowing primes equal to or bigger than $q$ to divide one of the $G_{i}(x, y)$ avoids generating duplicate relations, but misses relations having two large primes smaller than $q$ for $G_{j}(x, y)$ and a large prime larger than $q$ for $G_{3-j}(x, y)$. After we have sieved over all roots in $\mathcal{R}_{1}(q)$ and $\mathcal{R}_{2}(q)$ we take the next value of $b$; after we have sieved over all values of $b$ we take the next prime in the interval $\left[L^{(u)}, L^{(l)}\right]$. We implemented lattice sieving only for the case of two quadratic polynomials.

Since the sieving is the most time-consuming step of the algorithm, its implementation is critical. It is a lot of work to sieve over a small prime $p$, and just a small amount of $[\log p]$ is added to the array elements. Therefore we sieve only over primes and prime powers larger than 30 . Also we do not add
$[\log p]$ to all array elements for primes $p<B_{j}$ with $p \mid c_{j, d_{j}}$ and $p \mid b$, but we divide $c_{j}(a, b)$ by $p$. Furthermore we split the $a$-interval into subintervals that fit in the secondary cache of the computer, making the sieving faster. For a group of small primes, which consists of the primes for which we sieve over a power rather than over the prime itself, we again split the subintervals into smaller subintervals which fit into the primary cache. The user can install several "early abort" bounds: if the leftover part of $F_{i}(a, b)$ after trial division over all primes below a bound $B<B_{i}$ is bigger than a userspecified constant times the square of the large prime bound, then the pair $(a, b)$ is not considered to be a candidate for a relation and is thrown away. In the case of lattice sieving, the values of $a$ with $a \equiv b r_{1} r_{2}^{-1} \bmod q$ are far away from each other for a fixed value of $b$. In Section 5 we explained how we select polynomials such that we can increase the efficiency of the sieving by taking a huge $a$-interval and $b=1$. Therefore we call it line sieving.

## 7. THE FILTERING

The aim of filtering is just the reduction of the amount of data. We want to find a subset $\mathcal{S}$ of all relations $\{(a, b)\}$ found in the sieving step and a subset $\mathcal{T}$ of the set of free rational primes $p$ such that $\left(\prod_{\mathcal{T}} p\right)\left(\prod_{\delta}\left(a-b \alpha_{i}\right)\right)$ is a square in $\mathbb{Q}_{n}\left[\alpha_{i}\right]$, for $i=1,2$. Therefore every algebraic prime $p$ dividing one of the products $\left(\prod_{\mathcal{T}} p^{d_{i}}\right)\left(\prod_{\delta} F_{i}(a, b)\right)$ for a certain root $\left(r_{1}: r_{2}\right) \in \mathcal{R}_{i}(p)$ (from now on denoted by $\left.p_{\left(r_{1}: r_{2}\right)}\right)$ must occur to an even power with respect to this root. (Here we should see $p^{d_{i}}$ as the product of one factor $p$ for every root $\left.\left(r_{1}: r_{2}\right) \in \mathcal{R}_{i}(p)\right)$. A prime $p_{\left(r_{1}: r_{2}\right)}$ occurs in a relation $(a, b)$ for polynomial $i$ if $p$ divides $F_{i}(a, b)$ and $(a \bmod p: b \bmod p)=\left(r_{1}: r_{2}\right)$. A prime $p_{\left(r_{1}: r_{2}\right)}$ occurs in a free relation for both polynomials if $p$ is a free prime. We say that a prime $p_{\left(r_{1}: r_{2}\right)}$ occurs in a relation for polynomial $i$ if it occurs in a relation $(a, b)$ for polynomial $i$ or if $p$ is a free prime. It is obvious that a relation in which some prime $p_{\left(r_{1}: r_{2}\right)}$ occurs to an odd power for one of the two
polynomials is useless, if this prime is not occurring in some other relation to an odd power for the same polynomial. The filtering stage throws away such relations. If a prime $p_{\left(r_{1}: r_{2}\right)}$ occurs to an odd power in just two relations for the same polynomial and one of them belongs to the set $\mathcal{S}$, the other one should also be part of $\mathcal{S}$. In the filtering stage the two relations are grouped into a relationset. If one relation from a relation-set is chosen in the set $\mathcal{S}$, then all relations from that relationset should be in $\mathcal{S}$. By creating the relation-set we have eliminated the need to take care of the prime $p_{\left(r_{1}: r_{2}\right)}$ when looking for the set $\mathcal{S}$. In this way the amount of data and the size of the matrix for the next linear algebra step are reduced.

The relations found in the sieving step are read in sequentially. In order to regulate the amount of memory used, the user first chooses a number of temporary files among which the data will be distributed. During the filtering process data from only one temporary file will be in the working memory of the computer. A hash function is implemented that distributes the primes equally over the temporary files. For all the primes in the input file with norm larger than some user-determined bound $U \geq \max \left(W_{1}, W_{2}\right)$ and occurring to an odd power in one of the $F_{i}(a, b)$, the filter program calculates the index of the corresponding temporary file by using the hash function on the prime. The relation is written to the file with the smallest number it gets from all these primes. We store $a, b$, and the primes that were written in the input file. Extra features in this program have been added, such as looking only at the primes below some user-determined bound and throwing away all the relations containing a prime bigger than some user-determined bound.

When all relations are read in and stored in the corresponding files, the combining and throwing out process starts. First all relations $(a, b)$ of the first file are read in and stored in a heap [Standish 1980, § 3.7.1] in descending order according to the largest prime that led to the storage of the relation in this file. We start by considering the
relations in which the largest prime $p$ corresponding to this file occur. We calculate the root ( $r_{1}: r_{2}$ ) for this prime for the relations $(a, b)$ by calculating $(a \bmod p: b \bmod p)$. If $d_{1}+d_{2}$ different roots for this prime appear in these relations, we append the free relation for this prime. If, while looking at a prime $p_{\left(r_{1}: r_{2}\right)}$, we see the same relation $(a, b)$ twice, we throw out one of the occurrences. If some prime $p_{\left(r_{1}: r_{2}\right)}$ occurs exactly once for one of the polynomials, the corresponding relation is thrown out. If a prime $p_{\left(r_{1}: r_{2}\right)}$ occurs twice for one of the polynomials, the two corresponding relations are grouped in a relation-set. If the user wishes, for primes $p_{\left(r_{1}: r_{2}\right)}$ that occur just three times for one of the polynomials, the program replaces the three corresponding relations by two relation-sets of two relations each.

Next the resulting relation-sets and the relations that were not combined are stored at the next place corresponding to the hash function. If the relation or relation-set contains smaller primes to an odd power that correspond to the same file, we keep the relation(-set) in the working memory. We store the relation(-set) in the heap according to the largest of those primes. Otherwise, if the relation (-set) contains primes to an odd power that correspond to other files, we store the relation(-set) in the file among those with the smallest number exceeding the number of the file which we currently have in memory, or, if there is no such file, in the file among those with the smallest number. If the relation(-set) only contains larger primes to an odd power that correspond to the same file, we keep the relation(-set) in the same file, but write it to disk. If the relation(-set) does not contain any primes larger than $W_{i}$ to an odd power anymore, we write it to an output file. This circular queue is constructed in such a way that, when trying to throw out relation(-set)s or combining relation(-set)s into (new) relation-sets for some prime $p$, all relation (-set)s containing that prime to an odd power will be considered. When we store a relation-set we store all relations it contains, together with their primes exceeding $W_{i}$ and the free primes of the relation-set.

After the first temporary file is treated in this way, the same process takes place consecutively on the other files. Of course, relation-sets can also be combined with each other. Relation-sets that become too large, in the sense that they contain more relations than a user-determined bound, are thrown away, in an attempt to keep the matrix for the next linear algebra step sparse. The user can fix the maximum number of relation-sets that can be thrown away. When the last file has been processed the program starts again on the first file, until no changes have taken place in the last round or the number of passes has reached a user-chosen bound. Then all relation(-set)s that are still in the temporary files are written to the output file.

The filter program counts the number of relations and relation-sets that remain and the number of primes $p_{\left(r_{1}: r_{2}\right)}$ occurring to an odd power in one of the relations or relation-sets with norm larger than $U$. From these data one can estimate whether there are enough relations.
In practice we used filter several times for one number. To save disk space we chose big printing bounds, $W_{1}=W_{2}=10^{6}$, say. First we applied the filter program to the output of the sieving step with $U=W_{1}$. On the remaining relation(-set)s we applied the program factorrelations to compute the prime factors between a smaller bound $W^{\prime}$ and $W_{i}$ of $F_{1}(a, b)$ and $F_{2}(a, b)$ for all relations $(a, b)$ and store the relation(-set)s, now with all primes exceeding $W^{\prime}$. Then we again applied the filter program, now on all primes exceeding $U^{\prime}$, with $W^{\prime} \leq U^{\prime}<U$. These steps were repeated until we reached a bound below which many primes occur at least four times, so no combining or throwing out could be done, or until we were content with the resulting matrix size.

Another method that can be used for reducing the amount of data is structured Gaussian elimination, described in [LaMacchia et al. 1991], for example. A comparison between our filtering method and structured Gaussian elimination has not yet been made.

## 8. THE BLOCK LANCZOS METHOD

After enough relations have been collected and the filter program has reduced the amount of data, we try to find a subset $\mathcal{S}$ of the remaining relations and a subset $\mathcal{T}$ of the set of free primes such that $\left(\prod_{\mathcal{T}} p\right)\left(\prod_{\delta}\left(a-b \alpha_{i}\right)\right)$ is a square in $\mathbb{Q}_{n}\left[\alpha_{i}\right]$, for $i=1,2$. For simplicity, from now on we view a relation left after the filtering stage as a relation-set containing only one relation. To this collection of relation-sets we append a relation-set for every free prime below $\max \left(W_{1}, W_{2}\right)$, containing this prime.

A relation-set $\mathcal{V}$ consists of two (possibly empty) subsets $\mathcal{V}_{f}$ and $\mathcal{V}_{r}$ that contain the free primes and the relations of $\mathcal{V}$, respectively. For every relationset $\mathcal{V}$ we construct a vector

$$
\boldsymbol{v}(\mathcal{V})=\sum_{\mathcal{V}_{p}} \boldsymbol{v}(p)+\sum_{\mathcal{V}_{r}} \boldsymbol{v}(a, b) .
$$

The vectors $\boldsymbol{v}(a, b)$ are as described in Sections 2 and 3; the vectors $\boldsymbol{v}(p)$ are described in Section 4. We build a matrix $\boldsymbol{M}$ whose columns are all vectors $\boldsymbol{v}(\mathcal{V})$. We remove the rows that contain only zeros. They correspond to primes $\left(q,\left(r_{1}: r_{2}\right)\right)$ occurring to an even power in every relation-set $\mathcal{V}$. We want to calculate some nontrivial vectors of the null space of this matrix.

Since Gaussian elimination [Knuth 1981, § 4.6.2, Algorithm N ] requires too much memory for the large sparse matrices we have, we use a variation of the iterative Lanczos method. Proofs on both standard Lanczos and block Lanczos can be found in [Montgomery 1995]. The standard Lanczos algorithm starts with a symmetric, positive definite $k \times k$ matrix $\boldsymbol{A}$ over the field $K=\mathbb{R}$. If $\boldsymbol{b} \in \mathbb{R}^{k}$ we solve $\boldsymbol{A x}=\boldsymbol{b}$ by the following iterative procedure: set $\boldsymbol{w}_{0}=\boldsymbol{b}$ and

$$
\begin{equation*}
\boldsymbol{w}_{i}=\boldsymbol{A} \boldsymbol{w}_{i-1}-\sum_{j=0}^{i-1} c_{i j} \boldsymbol{w}_{j} \tag{8.1}
\end{equation*}
$$

for $i>0$, where

$$
c_{i j}=\frac{\boldsymbol{w}_{j}^{T} \boldsymbol{A}^{2} \boldsymbol{w}_{i-1}}{\boldsymbol{w}_{j}^{T} \boldsymbol{A} \boldsymbol{w}_{j}}
$$

It can be shown that after at most $k$ iterations we will find $\boldsymbol{w}_{i}=\mathbf{0}$. If $l$ is the first value of $i$ such that $\boldsymbol{w}_{i}=\mathbf{0}$, we have $\boldsymbol{w}_{i}^{T} \boldsymbol{A} \boldsymbol{w}_{i} \neq 0$ for $0 \leq i<l$, $\boldsymbol{w}_{j}^{T} \boldsymbol{A} \boldsymbol{w}_{i}=0$ for $i \neq j$, and $\boldsymbol{A} \mathcal{W} \subseteq \mathcal{W}$, where $\mathcal{W}$ is the span of $\boldsymbol{w}_{0}, \boldsymbol{w}_{1}, \ldots, \boldsymbol{w}_{l-1}$. One can deduce that

$$
\boldsymbol{x}=\sum_{i=0}^{l-1} \frac{\boldsymbol{w}_{i}^{T} \boldsymbol{b}}{\boldsymbol{w}_{i}^{T} \boldsymbol{A} \boldsymbol{w}_{i}} \boldsymbol{w}_{i}
$$

is a solution of $\boldsymbol{A} \boldsymbol{x}=\boldsymbol{b}$. Since $\boldsymbol{w}_{j}^{T} \boldsymbol{A}^{2} \boldsymbol{w}_{i-1}=0$ for $j<i-2$ we can simplify the calculation of $\boldsymbol{w}_{i}$ to

$$
\begin{equation*}
\boldsymbol{w}_{i}=\boldsymbol{A} \boldsymbol{w}_{i-1}-c_{i, i-1} \boldsymbol{w}_{i-1}-c_{i, i-2} \boldsymbol{w}_{i-2} \tag{8.2}
\end{equation*}
$$

for $i \geq 2$.
Standard Lanczos can also be applied over other fields, provided that $\boldsymbol{w}_{i}^{T} \boldsymbol{A} \boldsymbol{w}_{i} \neq 0$ when $\boldsymbol{w}_{i} \neq \mathbf{0}$ during this process. Working over the field $\mathbb{F}_{2}$ instead of $\mathbb{R}$ has the advantage that one can apply a matrix to $N$ different vectors simultaneously, where $N$ is the computer word size. Inspired by work of Coppersmith [1993], Montgomery [1995] implemented the block Lanczos method, which exploits this advantage.

The block Lanczos algorithm creates a sequence of subspaces $\mathcal{W}_{i}$ instead of vectors $\boldsymbol{w}_{i}$. Applying standard Lanczos over $\mathbb{F}_{2}$ has the problem that in approximately half of the cases the requirement $\boldsymbol{w}_{i}^{T} \boldsymbol{A} \boldsymbol{w}_{i} \neq 0$ if $\boldsymbol{w}_{i} \neq \mathbf{0}$ is violated. In the block Lanczos algorithm the analogous requirement is that no nonzero vector in $\mathcal{W}_{i}$ is $\boldsymbol{A}$-orthogonal to $\mathcal{W}_{i}$. This will hold when $\boldsymbol{W}_{i}^{T} \boldsymbol{A} \boldsymbol{W}_{i}$ is invertible, where $\boldsymbol{W}_{i}$ is a matrix whose column vectors span $\mathcal{W}_{i}$.

For $\boldsymbol{A}$ a symmetric $k \times k$ matrix over a field $K$ and $\boldsymbol{V}_{0}$ an arbitrary $k \times N$ matrix, the block Lanczos algorithm proceeds by setting, for $i=0,1, \ldots$,

$$
\begin{align*}
\boldsymbol{W}_{i} & =\boldsymbol{V}_{i} S_{i}  \tag{8.3}\\
\boldsymbol{C}_{i+1, j} & =\left(\boldsymbol{W}_{j}^{T} \boldsymbol{A} \boldsymbol{W}_{j}\right)^{-1} \boldsymbol{W}_{j}^{T} \boldsymbol{A}\left(\boldsymbol{A} \boldsymbol{W}_{i} \boldsymbol{S}_{i}^{T}+\boldsymbol{V}_{i}\right)  \tag{8.4}\\
\boldsymbol{V}_{i+1} & =\boldsymbol{A} \boldsymbol{W}_{i} \boldsymbol{S}_{i}^{T}+\boldsymbol{V}_{i}-\sum_{j=0}^{i} \boldsymbol{W}_{j} C_{i+1, j} \tag{8.5}
\end{align*}
$$

In (8.4), $j$ ranges from 0 through $i$. In (8.3) the $N \times N_{i}$ matrix $\boldsymbol{S}_{i}$ (where $N_{i} \leq N$ ) consists of zeros except for exactly one 1 per column and at most one 1 per row, thus selecting columns from $\boldsymbol{V}_{i}$ for $\boldsymbol{W}_{i}$. We choose the columns of $\boldsymbol{V}_{i}$ in such a way that the corresponding columns of $\boldsymbol{V}_{i}^{T} \boldsymbol{A} \boldsymbol{V}_{i}$ are a linearly independent spanning set of all columns of $\boldsymbol{V}_{i}^{T} \boldsymbol{A} \boldsymbol{V}_{i}$. Thus $N_{i}=\operatorname{rank}\left(\boldsymbol{V}_{i}^{T} \boldsymbol{A} \boldsymbol{V}_{i}\right)$, and one can prove that the resulting matrix $\boldsymbol{W}_{i}^{T} \boldsymbol{A} \boldsymbol{W}_{i}$ is invertible. The iteration process stops when $\boldsymbol{V}_{i}^{T} \boldsymbol{A} \boldsymbol{V}_{i}=$ $\mathbf{0}$, for $i=l$ say.

If $\boldsymbol{V}_{l}=\mathbf{0}$, the matrix $\boldsymbol{W}_{i}^{T} \boldsymbol{A} \boldsymbol{W}_{i}$ is invertible for $0 \leq i<l$, and we have

$$
\begin{align*}
\boldsymbol{W}_{j}^{T} \boldsymbol{A} \boldsymbol{W}_{i}=\mathbf{0} & \text { for } i \neq j  \tag{8.6}\\
\boldsymbol{W}_{j}^{T} \boldsymbol{A} \boldsymbol{V}_{i}=\mathbf{0} & \text { for } 0 \leq j<i \leq l \tag{8.7}
\end{align*}
$$

and finally $\boldsymbol{A} \mathcal{W} \subseteq \mathcal{W}$, where $\mathcal{W}$ is the span of $\mathcal{W}_{0}, \mathcal{W}_{1}, \ldots, \mathcal{W}_{l-1}$. If $\boldsymbol{b} \in \mathcal{W}$ one can further deduce that the vector

$$
\boldsymbol{x}=\sum_{i=0}^{l-1} \boldsymbol{W}_{i}\left(\boldsymbol{W}_{i}^{T} \boldsymbol{A} \boldsymbol{W}_{i}\right)^{-1} \boldsymbol{W}_{i}^{T} \boldsymbol{b}
$$

is a solution of $\boldsymbol{A} \boldsymbol{x}=\boldsymbol{b}$.
If we can choose $\boldsymbol{S}_{i}$ in such a way that $\operatorname{span}\left(\boldsymbol{V}_{i}\right) \subseteq$ $\operatorname{span}\left(\mathcal{W}_{0}, \mathcal{W}_{1}, \ldots, \mathcal{W}_{i+1}\right)$, it is possible to simplify the calculation of $\boldsymbol{V}_{i+1}$ in (8.5) in a way similar to that in which the calculation of $\boldsymbol{w}_{i}$ in (8.1) was simplified to (8.2):

$$
\begin{align*}
& \boldsymbol{V}_{i+1}=\boldsymbol{A} \boldsymbol{W}_{i} \boldsymbol{S}_{i}^{T}+\boldsymbol{V}_{i}-\boldsymbol{W}_{i} \boldsymbol{C}_{i+1, i} \\
&-\boldsymbol{W}_{i-1} \boldsymbol{C}_{i+1, i-1}-\boldsymbol{W}_{i-2} \boldsymbol{C}_{i+1, i-2} \tag{8.8}
\end{align*}
$$

for $i \geq 2$. The requirement is fulfilled when the columns of $\boldsymbol{V}_{i}$ are in $\operatorname{span}\left(\mathcal{W}_{i}, \mathcal{W}_{i+1}\right)$. From (8.8) we can deduce that

$$
\boldsymbol{V}_{i+1}=\boldsymbol{A} \boldsymbol{W}_{i} \boldsymbol{S}_{i}^{T}+\boldsymbol{V}_{i}-\boldsymbol{W}
$$

where $\boldsymbol{W}$ is a $k \times N$ matrix whose columns are linear combinations of the columns of $\boldsymbol{W}_{i}, \boldsymbol{W}_{i-1}$ and $\boldsymbol{W}_{i-2}$. Notice that the columns that were not selected from $\boldsymbol{V}_{i}$ are zero in $\boldsymbol{S}_{i}^{T}$ and in $\boldsymbol{A} \boldsymbol{W}_{i} \boldsymbol{S}_{i}^{T}$ as well. Therefore a nonselected column of $\boldsymbol{V}_{i}$ is equal to the sum of the corresponding columns of
$\boldsymbol{V}_{i+1}$ and $\boldsymbol{W}$. Using (8.7) and (8.6) we can deduce that such a column of $\boldsymbol{W}$ must be a linear combination of the columns of $\boldsymbol{W}_{i}$ only. Choosing the columns of $\boldsymbol{V}_{i+1}$ which were not selected in $\boldsymbol{V}_{i}$ guarantees that the nonselected columns of $\boldsymbol{V}_{i}$ are in $\operatorname{span}\left(\mathcal{W}_{i}, \mathcal{W}_{i+1}\right)$. If these columns of $\boldsymbol{V}_{i}$ are independent but the corresponding columns of $\boldsymbol{V}_{i+1}^{T} \boldsymbol{A} \boldsymbol{V}_{i+1}$ are dependent we cannot fulfill the requirement that $\boldsymbol{W}_{i+1}^{T} \boldsymbol{A} \boldsymbol{W}_{i+1}$ is invertible, and the algorithm fails. In practice this has never happened. Otherwise we choose a spanning set of columns for $\boldsymbol{V}_{i+1}^{T} \boldsymbol{A} \boldsymbol{V}_{i+1}$ including the columns that were not selected for $\boldsymbol{V}_{i}$, and choose $\boldsymbol{S}_{i+1}$ accordingly.

To apply the block Lanczos method to our matrix $\boldsymbol{M}$, we have to deal with several obstructions. First $\boldsymbol{M}$ need not be symmetric and therefore we apply the algorithm to the symmetric matrix $\boldsymbol{A}=$ $\boldsymbol{M}^{T} \boldsymbol{M}$. It is obvious that any solution of $\boldsymbol{M} \boldsymbol{x}=\mathbf{0}$ satisfies $\boldsymbol{A} \boldsymbol{x}=\mathbf{0}$, but the converse need not be true. Secondly, if we want to find a vector from the null space of $\boldsymbol{A}$ and start with $\boldsymbol{b}=\mathbf{0}$, we will find the trivial solution. We overcome this problem by starting with a random vector $\boldsymbol{y}$ and taking $\boldsymbol{b}=\boldsymbol{A} \boldsymbol{y}$. When $\boldsymbol{x}$ is a solution of $\boldsymbol{A x}=\boldsymbol{b}=\boldsymbol{A} \boldsymbol{y}$, then $\boldsymbol{x}-\boldsymbol{y}$ will be a random vector from the null space of $\boldsymbol{A}$. Thirdly, several vectors from the null space of $\boldsymbol{M}$ have to be found, since not every dependency corresponding with such a vector will lead to a nontrivial factor of $n$. Note that during the iteration steps the vector $\boldsymbol{b}$ only is involved in the calculation of the solution vector $\boldsymbol{x}$. Therefore replacing $\boldsymbol{b}$ by a $k \times N$ matrix $\boldsymbol{B}$ in the calculation of $\boldsymbol{x}$ will give a solution of $\boldsymbol{A} \boldsymbol{X}=\boldsymbol{B}$, with $\boldsymbol{X}$ also a $k \times N$ matrix. To find several vectors in the null space of $\boldsymbol{A}$ we start with a random $k \times N$ matrix $\boldsymbol{Y}$ and calculate solutions of $\boldsymbol{A} \boldsymbol{X}=\boldsymbol{A} \boldsymbol{Y}$. The $N$ column vectors of $\boldsymbol{X}-\boldsymbol{Y}$ will be random vectors in the null space of $\boldsymbol{A}$ and we extract the ones which are also in the null space of $\boldsymbol{M}$.

Final obstructions are the two requirements for $\boldsymbol{x}$ to be a solution of $\boldsymbol{A} \boldsymbol{x}=\boldsymbol{b}$. First $\boldsymbol{b}$ has to be in $\mathcal{W}=\operatorname{span}\left(\mathcal{W}_{0}, \mathcal{W}_{1}, \ldots, \mathcal{W}_{l-1}\right)$. This can be arranged by initializing $\boldsymbol{V}_{0}$ as $\boldsymbol{A} \boldsymbol{Y}$, where $\boldsymbol{Y}$ is a
random $k \times N$ matrix. Secondly, the algorithm often terminates with $\boldsymbol{V}_{l}^{T} \boldsymbol{A} \boldsymbol{V}_{l}=\mathbf{0}$ but $\boldsymbol{V}_{l} \neq \mathbf{0}$. Montgomery presumes that the column vectors of $\boldsymbol{A}(\boldsymbol{X}-\boldsymbol{Y})$ and $\boldsymbol{A} \boldsymbol{V}_{l}$ are both in $\operatorname{span}\left(\boldsymbol{V}_{l}\right)$, which has maximal rank $N$, but in practice the rank is much smaller. We may expect that some linear combinations of these vectors are in the null space of $\boldsymbol{A}$. Combined with the need to find vectors in the null space of $\boldsymbol{M}$ instead of $\boldsymbol{A}$, it suffices to construct a suitable matrix $\boldsymbol{U}$ such that $\boldsymbol{M} \boldsymbol{Z} \boldsymbol{U}=$ $\mathbf{0}$, where $\boldsymbol{Z}$ is a $k \times 2 N$ matrix of the columns of $\boldsymbol{X}-\boldsymbol{Y}$ and $\boldsymbol{V}_{l}$. We first compute $\boldsymbol{M} \boldsymbol{Z}$. Then we determine a matrix $\boldsymbol{U}$ whose columns span the null space of $\boldsymbol{M} \boldsymbol{Z}$. The output is a basis for $\boldsymbol{Z} \boldsymbol{U}$.

For implementing the calculation of $\boldsymbol{V}_{i+1}$ in (8.8) one can bring further down the number of calculations by using the following steps: for $i=0,1, \ldots$, set
$\boldsymbol{V}_{i+1}=\boldsymbol{A} \boldsymbol{V}_{i} \boldsymbol{S}_{i} \boldsymbol{S}_{i}^{T}+\boldsymbol{V}_{i} \boldsymbol{D}_{i+1}+\boldsymbol{V}_{i-1} \boldsymbol{E}_{i+1}+\boldsymbol{V}_{i-2} \boldsymbol{F}_{i+1}$,
where

$$
\begin{aligned}
\boldsymbol{D}_{i+1}= & \boldsymbol{I}_{N}-\boldsymbol{W}_{i}^{*}\left(\boldsymbol{V}_{i}^{T} \boldsymbol{A}^{2} \boldsymbol{V}_{i} \boldsymbol{S}_{i} \boldsymbol{S}_{i}^{T}+\boldsymbol{V}_{i}^{T} \boldsymbol{A} \boldsymbol{V}_{i}\right), \\
\boldsymbol{W}_{i}^{*}= & \boldsymbol{S}_{i}\left(\boldsymbol{S}_{i}^{T} \boldsymbol{V}_{i}^{T} \boldsymbol{A} \boldsymbol{V}_{i} \boldsymbol{S}_{i}\right)^{-1} \boldsymbol{S}_{i}^{T} \\
\boldsymbol{E}_{i+1}= & -\boldsymbol{W}_{i-1}^{*} \boldsymbol{V}_{i}^{T} \boldsymbol{A} \boldsymbol{V}_{i} \boldsymbol{S}_{i} \boldsymbol{S}_{i}^{T} \\
\boldsymbol{F}_{i+1}= & -\boldsymbol{W}_{i-2}^{*}\left(\boldsymbol{I}_{N}-\boldsymbol{V}_{i-1}^{T} \boldsymbol{A} \boldsymbol{V}_{i-1} \boldsymbol{W}_{i-1}^{*}\right) \\
& \left(\boldsymbol{V}_{i-1}^{T} \boldsymbol{A}^{2} \boldsymbol{V}_{i-1} \boldsymbol{S}_{i-1} \boldsymbol{S}_{i-1}^{T}+\boldsymbol{V}_{i-1}^{T} \boldsymbol{A} \boldsymbol{V}_{i-1}\right) \boldsymbol{S}_{i} \boldsymbol{S}_{i}^{T},
\end{aligned}
$$

and where, for $i<0, \boldsymbol{W}_{i}^{*}$ and $\boldsymbol{V}_{i}$ are $\mathbf{0}$ and $\boldsymbol{S}_{i}$ is $\boldsymbol{I}_{N}$. When $\boldsymbol{S}_{i-1}=\boldsymbol{I}_{N}$ then $\boldsymbol{F}_{i+1}=\mathbf{0}$ and the term $\boldsymbol{V}_{i-2} \boldsymbol{F}_{i+1}$ in the expression for $\boldsymbol{V}_{i+1}$ can be omitted.

For large sparse matrices Lanczos' algorithm requires less storage than Gaussian elimination. It only needs the original matrix and some extra vectors of length $k$ and some $N \times N$ matrices, while Gaussian elimination causes fill-in and therefore needs approximately $k^{2}$ bits. When $M$ has $d$ nonzero entries per row on average, the time needed by block Lanczos is $\mathcal{O}\left(d k^{2} / N\right)+\mathcal{O}\left(k^{2}\right)$. When $d$ is much smaller than $k$ this is considerably better than $\mathcal{O}\left(k^{3} / N\right)$ for Gaussian elimination.

In practice we make $\boldsymbol{M}$ extra sparse by removing the first row containing only ones and not append-
ing any character rows. Also one could implement the possibility to remove some of the dense rows corresponding to small primes. If $\boldsymbol{M}$ is a $k_{1} \times k_{2}$ matrix, the output of the block Lanczos algorithm will consist of a $k_{2} \times N$ matrix $\boldsymbol{P}$ with $N$ "pseudodependencies" of which we still have to find linear combinations to get a set $\mathcal{S}$ we look for. We solve this problem here, although in our implementation it is a part of the square root program. For several quadratic characters $\left(q,\left(s_{1}: s_{2}\right)\right)$ - chosen as described in Section 2 with $q$ larger than any prime dividing any $F_{i}(a, b)$ - we form a vector $\boldsymbol{q}_{\left(q,\left(s_{1}: s_{2}\right)\right)}$ of length $k_{2}$ by inserting a zero for all of the $k_{2}$ relation-sets $\mathcal{V}$ with

$$
\prod_{v_{p}}\left(\frac{p}{q}\right) \prod_{\mathcal{V}_{f}}\left(\frac{a-b\left(s_{1} s_{2}^{-1} \bmod q\right)}{q}\right)=1
$$

and a one otherwise. A vector of length $N$ orthogonal to all vectors $\boldsymbol{q}_{\left(q,\left(s_{1}: s_{2}\right)\right)} \boldsymbol{P}$ is indicating a linear combination of the $N$ pseudo-dependencies which is favourable to all chosen quadratic characters. We construct a basis for the space orthogonal to all vectors $\boldsymbol{q}_{\left(q,\left(s_{1}: s_{2}\right)\right)} \boldsymbol{P}$. Each of these basis vectors indicates which pseudo-dependencies of $\boldsymbol{P}$ should be combined for a real dependency, thereby indicating a set $\delta$.

## 9. EXTRACTING THE SQUARE ROOT

At this stage we have two squares $\beta^{2}=\left(\prod_{p \in \mathcal{T}} p\right) \times$ $\left(\prod_{(a, b) \in \mathcal{S}}\left(a-b \alpha_{1}\right)\right)$ and $\gamma^{2}=\left(\prod_{p \in \mathcal{T}} p\right)\left(\prod_{(a, b) \in \mathcal{S}}(a-\right.$ $\left.\left.b \alpha_{2}\right)\right)$ in $\mathbb{Q}_{n}\left[\alpha_{1}\right]$ and $\mathbb{Q}_{n}\left[\alpha_{2}\right]$, respectively. We have to calculate $\beta$ and $\gamma$. If we write both squares as polynomials of degree less than $d_{i}$ in $\alpha_{i}$, the coefficients will be gigantic. Then a conventional method such as the one described in [Cohen 1993, §3.6.2] cannot be used. Couveignes [1993] calculates the square roots modulo several primes and applies the Chinese Remainder Theorem, a method that presently works only for number fields of odd degree.

Montgomery [1994] attacks the problem using an iterative process. He starts by partitioning the set $\mathcal{S}$ in two subsets $\mathcal{S}_{1}$ and $\mathcal{S}_{2}$ and the set $\mathcal{T}$ in two
subsets $\mathfrak{T}_{1}$ and $\mathcal{T}_{2}$ to advance the cancellation of primes $p_{\left(r_{1}: r_{2}\right)}$ in both products

$$
\begin{equation*}
\frac{\prod_{p \in \mathcal{J}_{1}} p^{d_{i}} \prod_{(a, b) \in \mathcal{S}_{1}} F_{i}(a, b)}{\prod_{p \in \mathcal{T}_{2}} p^{d_{i}} \prod_{(a, b) \in \mathcal{S}_{2}} F_{i}(a, b)} \tag{9.1}
\end{equation*}
$$

$i=1,2$. Here, the expressions $p^{d_{i}}$ for the free primes $p \in\left\{\mathcal{T}_{1} \cup \mathcal{T}_{2}\right\}$ should be seen as the product of one factor $p$ for every root $\left(r_{1}: r_{2}\right) \in \mathcal{R}_{i}(p)$. We can for example choose $\mathcal{S}_{1}=\mathcal{S}, \mathcal{S}_{2}$ empty, $\mathcal{T}_{1}=\mathcal{T}$ and $\mathfrak{T}_{2}$ empty; or we can distribute $\mathcal{S}$ and $\mathfrak{T}$ over the sets $\mathcal{S}_{j}$ and $\mathcal{T}_{j}$ randomly. At the end of this section we will see how we tried to optimize this selection.

Set

$$
\begin{aligned}
& \nu_{1}^{2}=\frac{\prod_{\mathcal{J}_{1}} p \prod_{\delta_{1}}\left(a-b \alpha_{1}\right)}{\prod_{\mathcal{I}_{2}} p \prod_{\delta_{2}}\left(a-b \alpha_{1}\right)}, \\
& \nu_{2}^{2}=\frac{\prod_{\mathcal{J}_{1}} p \prod_{\delta_{1}}\left(a-b \alpha_{2}\right)}{\prod_{\mathcal{J}_{2}} p \prod_{\delta_{2}}\left(a-b \alpha_{2}\right)}
\end{aligned}
$$

then we will calculate $\nu_{1}$ and $\nu_{2}$, for which the congruence $\left(\varphi_{1}\left(\nu_{1}\right)\right)^{2} \equiv\left(\varphi_{2}\left(\nu_{2}\right)\right)^{2} \bmod n$ holds. The following algorithm is applied twice, first to calculate $\nu=\nu_{1}$ and then to calculate $\nu=\nu_{2}$. In the rest of this section we suppress the index $i$ when referring to $\nu_{i}, f_{i}, d_{i}, c_{i, k}$ and $\alpha_{i}$.

Starting with $\tau_{1}=\nu^{2}$, where $\nu$ is unknown, we will approximate in iteration step $j \geq 1$ the numerator (if $j$ is odd) or the denominator (if $j$ is even) of $\sqrt{\tau_{j}}$ by $\eta_{j}$ (to be explained below) and calculate $\tau_{j+1}$ using the formula

$$
\begin{equation*}
\tau_{j+1}=\tau_{j} \cdot\left(\eta_{j}^{2}\right)^{(-1)^{j}} \tag{9.2}
\end{equation*}
$$

Hence

$$
\begin{equation*}
\nu^{2}=\tau_{j+1} \frac{\prod_{l=1}^{\left\lfloor\frac{i+1}{2}\right\rfloor} \eta_{2 l-1}^{2}}{\prod_{l=1}^{\left\lfloor\frac{\left.L_{2}\right\rfloor}{2}\right\rfloor} \eta_{2 l}^{2}} . \tag{9.3}
\end{equation*}
$$

The product of the norms of the numerator and the denominator of $\tau_{j+1}$ in (9.3) will decrease at every iteration step. Small norms of numerator and denominator, however, do not guarantee that the coefficients of $\tau_{j+1}$ as a polynomial of degree $\leq d-1$ in $\alpha$ are small. Let $\alpha_{1}, \alpha_{2}, \ldots, \alpha_{d}$ be the conjugates of $\alpha$. For any polynomial $h(x) \in \mathbb{Q}[x]$
of degree at most $d-1$ the Lagrange interpolation formula gives

$$
h(x)=\sum_{k=0}^{d-1} x^{k} \sum_{l=1}^{d} h\left(\alpha_{l}\right) c_{k l}
$$

where the $c_{k l}$ can be calculated from

$$
\frac{f(x)}{\left(x-\alpha_{l}\right) f^{\prime}\left(\alpha_{l}\right)}=\sum_{k=0}^{d-1} c_{k l} x^{k}
$$

Therefore we can bound the coefficients of $h(x)$ in terms of the $\left|h\left(\alpha_{l}\right)\right|$. We use this observation by choosing the approximation $\eta_{j}$ in such a way that not only the product of the numerator and the denominator norms of the successive $\tau_{j+1}$ 's decreases, but also $\left|\tau_{j+1}\left(\alpha_{l}\right)\right|$ tends to decrease for all $l$. When the norms and the embeddings become small enough, we will express the final $\tau_{j+1}$ as a polynomial of degree $\leq d-1$ in $\alpha$ and find its square root by using the computer package PARI [Batut et al. 1995].

In order to find the $\eta_{j}$ 's we work with ideals. Denote by $\mathcal{O}$ the ring of integers in $\mathbb{Q}[\alpha]$ and for $x_{j} \in \mathbb{Q}[\alpha]$ by $\left\langle x_{1}, \ldots, x_{n}\right\rangle \mathcal{O}$ the fractional ideal generated by $x_{1}, \ldots, x_{n}$ in $\mathcal{O}$. Suppose

$$
\begin{equation*}
\left\langle\sqrt{\tau_{j}}\right\rangle \mathcal{O}=\frac{\prod_{l \in \mathcal{L}_{j 1}} \mathcal{P}_{l}^{c_{l}}}{\prod_{l \in \mathcal{L}_{j 2}} \mathcal{P}_{l}^{c_{l}}} \tag{9.4}
\end{equation*}
$$

is the factorization of $\left\langle\sqrt{\tau_{j}}\right\rangle \mathcal{O}$ into prime ideals $\mathcal{P}_{l}$ of $\mathcal{O}$, where $c_{l} \in \mathbb{Z}_{+}$for all $l$. At each iteration step we select an ideal $\mathcal{J}$ dividing the numerator (if $j$ is odd) or the denominator (if $j$ is even) of (9.4). The approximation $\eta_{j}$ will be a 'small' element of $\mathcal{J}$. If $\mathcal{J}$ divides the numerator, we divide $\left\langle\sqrt{\tau_{j}}\right\rangle \mathcal{O}$ by $\left\langle\eta_{j}\right\rangle \mathcal{O}$ and this will result in the disappearance of $\mathcal{J}$ in the numerator of $\left\langle\sqrt{\tau_{j}}\right\rangle \mathcal{O}$ and the appearance of a new integral ideal $\mathcal{Q}$ of $\mathcal{O}$ in the denominator of $\left\langle\sqrt{\tau_{j}}\right\rangle \mathcal{O}$. If $\mathcal{J}$ divides the denominator the converse happens. If $N(\mathcal{J})$ is sufficiently large, than $N(Q)$ will be much smaller than $N(\mathcal{J})$. In this way the product of the numerator and the denominator norms will decrease every iteration step.

To factor $\left\langle\tau_{j}\right\rangle \mathcal{O}$ into prime ideals, we use the ideal

$$
\begin{aligned}
& \mathcal{J}=\left\langle c_{d}, c_{d} \alpha+c_{d-1}, c_{d} \alpha^{2}+c_{d-1} \alpha+c_{d-2}, \ldots\right. \\
& \left.c_{d} \alpha^{d-1}+c_{d-1} \alpha^{d-2}+\cdots+c_{1}\right\rangle \mathcal{O}
\end{aligned}
$$

We have $\mathcal{J} \subseteq \mathcal{O}$ and $\langle 1, \alpha\rangle \mathcal{O J}=\mathcal{O}$ [Montgomery 1994]. From this we deduce $\langle a-b \alpha\rangle \mathcal{O} \subseteq \mathcal{O}$. Therefore, if we multiply an ideal $\langle a-b \alpha\rangle \mathcal{O}$ by $\mathcal{J}$ we can factor the result in prime ideals, all with a positive exponent. The ideals $\langle p\rangle \mathcal{O}$ with $p \in \mathcal{T}$ are already integral, so multiplication with $\mathcal{J}$ is not necessary for these ideals. We start with

$$
\begin{equation*}
\left\langle\tau_{1}\right\rangle \mathcal{O}=\frac{\mathfrak{f}^{\# \delta_{2}} \prod_{\mathcal{T}_{1}}\langle p\rangle \mathcal{O} \prod_{\mathcal{S}_{1}}\{\langle a-b \alpha\rangle \mathcal{O}\}}{\mathfrak{J \# \delta _ { 1 }} \prod_{\mathcal{T}_{2}}\langle p\rangle \mathcal{O} \prod_{\mathcal{S}_{2}}\{\langle a-b \alpha\rangle \mathcal{O}\}} \tag{9.5}
\end{equation*}
$$

and factor the ideals $\langle p\rangle \mathcal{O}$ and $\langle a-b \alpha\rangle \mathcal{O}$ into prime ideals. Therefore we split the primes in two subsets: the set of "special" primes which divide the index $[\mathcal{O}: \mathbb{Z}[\alpha]]$ and the remaining primes which we call normal. To every prime $p$ and every root $\left(r_{1}: r_{2}\right) \in \mathcal{R}(p)$ there correspond prime ideals dividing $\langle p\rangle \mathcal{O}$ if $p$ is an element of $\mathcal{T}_{1} \cup \mathcal{T}_{2}$ or dividing $\langle a-b \alpha\rangle \mathcal{O}$ if $p$ divides $F(a, b)\left((a, b) \in \mathcal{S}_{1} \cup \mathcal{S}_{2}\right)$. For a special prime there may exist more prime ideals corresponding to the same root, but for a normal prime $p$ the prime ideal $\mathcal{P}$ corresponding to a root $\left(r_{1}: r_{2}\right) \in \mathcal{R}(p)$ is uniquely determined. Based on practical experience Montgomery suspects - which we cannot prove - that in the latter case the correspondence is given by

$$
\mathcal{P}= \begin{cases}\left\langle p, c_{d} \alpha-c_{d} r_{1} r_{2}^{-1} \bmod p\right\rangle \mathcal{O} & \text { if } p \nmid c_{d},  \tag{9.6}\\ \langle p\rangle \mathcal{O}+\mathcal{J} & \text { if } p \mid c_{d}, r_{2}=0 \\ \mathfrak{J} \cdot\left\langle p, \alpha-r_{1} r_{2}^{-1} \bmod p\right\rangle \mathcal{O} & \text { if } p \mid c_{d}, r_{2} \neq 0\end{cases}
$$

For the special primes $p$ and for all their roots $\left(r_{1}: r_{2}\right) \in \mathcal{R}(p)$, we calculate the ideal $\mathcal{P}$ using (9.6) and factor it into prime ideals with help of the computer package PARI. While we read in the free primes and relations we accumulate a product of the factors of the right hand side of (9.5). We make a hash table containing an entry for each normal prime $p$ and $\left(r_{1}: r_{2}\right) \in \mathcal{R}(p)$ we encounter. Each entry contains the exponent of the corresponding prime ideal in the accumulated product so far: if
we meet a normal prime $p_{\left(r_{1}: r_{2}\right)}$ dividing a free prime or $F(a, b)$ in the numerator (or denominator) of $\left\langle\tau_{1}\right\rangle \mathcal{O}$ to the power $x$, we add (or subtract) $x$ to this exponent. If we encounter a special prime $p_{\left(r_{1}: r_{2}\right)}$, dividing a free prime $p$ or one of the $F(a, b)$, we use PARI to calculate the valuation of $\langle p\rangle \mathcal{O}$ or $\langle a-b \alpha\rangle \mathcal{O}$, respectively, at the ideals of $p_{\left(r_{1}: r_{2}\right)}$ we computed earlier. Also for these special ideals we keep track of the exponent of the ideal in the accumulated product so far. We also have to keep track of the exponent of $\mathcal{J}$ in the accumulated product. For each ideal $\langle a-b \alpha\rangle \mathcal{O}$ in the numerator or denominator we add or subtract 1 , respectively. When we have read in all free primes and relations we have factored $\left\langle\tau_{1}\right\rangle \mathcal{O}$ into prime ideals and a power of $\mathcal{J}$.

Now we start the iterative approximation process in which we use the LLL lattice basis reduction algorithm [Lenstra et al. 1982]. Assume we want to simplify the numerator. The algorithm selects an ideal $\mathcal{J}=\mathcal{P}_{1}^{s_{1}} \ldots \mathcal{P}_{k}^{s_{k}}$ dividing the numerator of $\left\langle\sqrt{\tau_{j}}\right\rangle \mathcal{O}$, with $s_{r}>0$ for all $r$. $N(\mathcal{J})$ should be chosen as large as computationally convenient for this first lattice basis reduction. Let $\left\{a_{1}, \ldots, a_{d}\right\}$ be an integral bases of $\mathcal{O}$. With help of PARI we construct a basis in Hermite Normal Form (HNF) [Cohen 1993, §4.7.2] expressed in $\left\{a_{1}, \ldots, a_{d}\right\}$ for each prime ideal $\mathcal{P}_{l}$ occurring in $\mathcal{J}$. PARI uses these bases to construct a basis of $\mathcal{J}$, also in HNF expressed in $\left\{a_{1}, \ldots, a_{d}\right\}$. Then we apply a lattice basis reduction to these $d$ basis vectors of $\mathcal{J}$. We find a basis of $\mathcal{J}$ consisting of 'small' vectors. In practice, when using one of these small vectors for our approximation $\eta_{j}$, the norm of the numerator of $\left\langle\tau_{j+1}\right\rangle \mathcal{O}$ will decrease by a factor $N(\mathcal{J})$ in comparison with the norm of the numerator of $\left\langle\tau_{j}\right\rangle \mathcal{O}$. In comparison with the norm of the denominator of $\left\langle\tau_{j}\right\rangle \mathcal{O}$, the norm of the denominator of $\left\langle\tau_{j+1}\right\rangle \mathcal{O}$ will increase by a factor much smaller than $N(\mathcal{J})$.

We apply a second lattice basis reduction to a slight modification of the basis which we find after the first lattice basis reduction, to search for an element $\eta_{j}$ in $\mathcal{J}$, which still has the same effect on the norm of $\left\langle\tau_{j+1}\right\rangle \mathcal{O}$, but yields small $\left|\tau_{j+1}\left(\alpha_{l}\right)\right|$ for
all $l$. Let $\boldsymbol{v}^{(1)}, \boldsymbol{v}^{(2)}, \ldots, \boldsymbol{v}^{(d)}$ be the reduced basis after the first lattice basis reduction, where $\boldsymbol{v}^{(r)}=$ $\sum_{k=0}^{d-1} v_{k}^{(r)} \alpha^{k}$. While we read in the free primes and relations we calculate an approximation of $\tau_{1}\left(\alpha_{l}\right)$ for $1 \leq l \leq d$. We choose $c$ such that

$$
c^{d}=\frac{L_{\max }\left(N\left(\tau_{j}\right)\right)^{1 / 2}}{N(\mathcal{J})\left(\operatorname{Disc}\left(f / c_{d}\right)\right)^{1 / 2}}
$$

where $\operatorname{Disc}(g)$ is the discriminant of the polynomial $g$ and $L_{\max }=10^{100}$. If all conjugates $\alpha_{1}, \ldots, \alpha_{d}$ of $\alpha$ are real, we construct the vectors

$$
\begin{aligned}
T \boldsymbol{v}^{(r)}=\left(v_{0}^{(r)}, v_{1}^{(r)},\right. & \ldots, v_{d-1}^{(r)}, \\
& \left.c \frac{\boldsymbol{v}^{(r)}\left(\alpha_{1}\right)}{\left|\tau_{j}\left(\alpha_{1}\right)\right|^{1 / 2}}, \ldots, c \frac{\boldsymbol{v}^{(r)}\left(\alpha_{d}\right)}{\left|\tau_{j}\left(\alpha_{d}\right)\right|^{1 / 2}}\right)^{T}
\end{aligned}
$$

for $1 \leq r \leq d$. If $\alpha_{s}$ and $\alpha_{t}$ are complex conjugates, we replace

$$
c \frac{\boldsymbol{v}^{(r)}\left(\alpha_{s}\right)}{\left|\tau_{j}\left(\alpha_{s}\right)\right|^{1 / 2}} \quad \text { by } \quad c \sqrt{2} \frac{\operatorname{Re}\left(\boldsymbol{v}^{(r)}\left(\alpha_{s}\right)\right)}{\left|\tau_{j}\left(\alpha_{s}\right)\right|^{1 / 2}}
$$

and

$$
c \frac{\boldsymbol{v}^{(r)}\left(\alpha_{t}\right)}{\left|\tau_{j}\left(\alpha_{t}\right)\right|^{1 / 2}} \quad \text { by } \quad c \sqrt{2} \frac{\operatorname{Im}\left(\boldsymbol{v}^{(r)}\left(\alpha_{s}\right)\right)}{\left|\tau_{j}\left(\alpha_{s}\right)\right|^{1 / 2}}
$$

in the construction of $T \boldsymbol{v}^{(r)}$. In this way all entries of the $T \boldsymbol{v}^{(r)}$ will be real and the absolute value of the determinant of the matrix formed by the last $d$ entries of these vectors remains the same. The determinant equals $\pm L_{\text {max }}$, which constant has been chosen in such a way that the second lattice basis reduction algorithm performs well. We apply the second lattice basis reduction to the vectors $\left\{T \boldsymbol{v}^{(r)}\right\}_{r=1}^{d}$ and take the first $d$ coordinates of one of the resulting vectors for $\eta_{j}$. When dividing $\tau_{j}$ by $\eta_{j}^{2}$ the ideal $\mathcal{J}$ in the numerator of $\left\langle\tau_{j}\right\rangle \mathcal{O}$ will disappear. At the same time the denominator of $\left\langle\tau_{j}\right\rangle \mathcal{O}$ will be multiplied by the square of a new ideal

$$
\begin{equation*}
\left\langle\eta_{j}\right\rangle \mathcal{O} / \mathcal{J}=: Q \tag{9.7}
\end{equation*}
$$

In practice also for this $\eta_{j}$ we have that $N(\mathbb{Q})$ is much smaller than $N(\mathcal{J})$.

If we simplify the denominator of $\tau_{j}$ we select an ideal $\mathcal{J}=\mathcal{P}_{1}^{s_{1}} \ldots \mathcal{P}_{k}^{s_{k}}$ which divides the denominator of $\left\langle\sqrt{\tau_{j}}\right\rangle \mathcal{O}$. For the first LLL reduction the algorithm proceeds as above. For the second LLL reduction the vectors $T \boldsymbol{v}^{(r)}$ become

$$
\begin{aligned}
& T \boldsymbol{v}^{(r)}=\left(v_{0}^{(r)}, \ldots, v_{d-1}^{(r)},\right. \\
& \left.\quad c \boldsymbol{v}^{(r)}\left(\alpha_{1}\right)\left|\tau_{j}\left(\alpha_{1}\right)\right|^{1 / 2}, \ldots, c \boldsymbol{v}^{(r)}\left(\alpha_{d}\right)\left|\tau_{j}\left(\alpha_{d}\right)\right|^{1 / 2}\right)^{T}
\end{aligned}
$$

for $1 \leq r \leq d$, and we replace $c \boldsymbol{v}^{(r)}\left(\alpha_{s}\right)\left|\tau_{j}\left(\alpha_{s}\right)\right|^{1 / 2}$ by $c \sqrt{2} \operatorname{Re}\left(\boldsymbol{v}^{(r)}\left(\alpha_{s}\right)\right)\left|\tau_{j}\left(\alpha_{s}\right)\right|^{1 / 2}$ and
$c \boldsymbol{v}^{(r)}\left(\alpha_{t}\right)\left|\tau_{j}\left(\alpha_{t}\right)\right|^{1 / 2}$ by $c \sqrt{2} \operatorname{Im}\left(\boldsymbol{v}^{(r)}\left(\alpha_{s}\right)\right)\left|\tau_{j}\left(\alpha_{s}\right)\right|^{1 / 2}$
in the construction of $T \boldsymbol{v}^{(r)}$ if $\alpha_{s}=\bar{\alpha}_{t}$. The constant $c$ should be calculated from

$$
c^{d}=\frac{L_{\max }}{N(\mathcal{J})\left(\operatorname{Disc}\left(f / c_{d}\right)\right)^{1 / 2}\left(N\left(\tau_{j}\right)\right)^{1 / 2}}
$$

In the next iteration step we avoid factoring $Q$ into prime ideals by including this ideal as a factor of the new $\mathcal{J}$. Its basis in HNF is found by using (9.7). Furthermore we need the embeddings of $\tau_{j}$ for the second LLL reduction. Using (9.2) we can calculate $\left|\tau_{j+1}\left(\alpha_{l}\right)\right|$ from $\left|\tau_{j}\left(\alpha_{l}\right)\right|$ and $\left|\eta_{j}\left(\alpha_{l}\right)\right|^{2}$. We stop with the iterative process when the norms of numerator and denominator and the embeddings of $\tau_{j+1}$ are small enough.

Next we calculate the square root $\sqrt{\tau_{j+1}}$ with help of PARI. We first write $\tau_{j+1}$ as a polynomial in $\alpha$. We construct an integer $t$, being the product of the index and the norms of all ideals which are still in the denominator of $\tau_{j+1}$. Hence $t \tau_{j+1}$ is a polynomial of degree $d$ in $\alpha$ with coefficients in $\mathbb{Z}$. During the algorithm we keep track of the coefficients of the numerator and the denominator of $\tau_{j}$ as a polynomial in $\alpha$ of degree $<d$ modulo several large primes. We use this to express the final $t \tau_{j+1}$ as a polynomial in $\alpha$ of degree $<d$ modulo these primes and we use the Chinese Remainder Theorem to find its coefficients in $\mathbb{Z}$. We divide this element of $\mathbb{Z}[\alpha]$ by $t$ and find its square root by using the method mentioned in [Cohen 1993, § 3.6.2]. This completes the computation of $\nu$.

We now apply the homomorphism $\varphi_{i}$, for $i=$ 1,2 , to the expressions found for $\nu_{i}$ :

$$
\varphi_{i}\left(\nu_{i}\right)=\sqrt{\tau_{j+1}}(m) \frac{\prod_{l=1}^{\left\lfloor\frac{j+1}{2}\right\rfloor} \eta_{2 l-1}(m)}{\prod_{l=1}^{\left\lfloor\frac{i}{2}\right\rfloor} \eta_{2 l}(m)} \bmod n .
$$

Here the $\eta_{i}(m) \bmod n$ are calculated and multiplied with the $\eta_{j}(m) \bmod n$, for $j<i$, straight after $\eta_{i}$ has been calculated, so there is no need to store a history. We calculate $\operatorname{gcd}\left(\varphi_{1}\left(\nu_{1}\right)-\varphi_{2}\left(\nu_{2}\right), n\right)$ and hope to find a nontrivial factor of $n$.

In practice the second lattice basis reduction applied to the $T \boldsymbol{v}^{(r)}$ will yield linear combinations of the $\boldsymbol{v}^{(1)}\left(\alpha_{i}\right), \ldots, \boldsymbol{v}^{\left(d_{i}\right)}\left(\alpha_{i}\right)$ with small coefficients. Therefore we can round the entries of the $T \boldsymbol{v}^{(r)}$ without introducing a lot of round-off accumulation. This is the reason why we do not perform one single lattice basis reduction. Now both reductions use integer arithmetic.

It is important for the speed of the algorithm to select the sets $\mathcal{S}_{1}, \mathcal{S}_{2}, \mathfrak{T}_{1}$ and $\mathfrak{T}_{2}$ such that we get as much cancellation of primes $p_{\left(r_{1}: r_{2}\right)}$ as possible. We start with putting half the number of relations of $\mathcal{S}$ and half the number of free primes of $\mathcal{T}$ in $\mathcal{S}_{1}$ and $\mathcal{T}_{1}$, respectively, and the rest in $\mathcal{S}_{2}$ and $\mathcal{T}_{2}$. While we read in all relations and free primes for one of the polynomials, $f_{i}(x)$ say, and accumulate the prime ideal factorization of the numerator and the denominator of $\nu_{i}$, we decide whether it is profitable to put the current relation $(a, b)$ or free prime $p$ in the denominator while it was originally scheduled for the numerator (or vice-versa). If we decide to do so, then we put this relation for this $f_{i}$ in the denominator and compensate this by multiplying the final $\varphi_{i}\left(\nu_{i}\right)$ with $a-b m \bmod n$ or $p \bmod n$ respectively.

When using PARI for calculations in number fields it is necessary to use the function initalg, which calculates amongst others an integral basis. This function needs to factor the discriminant of the polynomial, which can be too hard for PARI. We solve this problem by factoring the discriminant ourselves and giving the primes to PARI with the function addprimes.

## 10. EXPERIMENTAL RESULTS

In this section we summarize factorization runs for several integers of up to 162 digits, indicating the time spent on each major step of the algorithm.

Except for $72^{99}+1$, all numbers were initially exposed to trial division and the elliptic curve method [Cohen 1993, §10.3] to find the factors below 40 digits. Thus, in the tables, "C98 from $7^{128}+6^{128}$ " means a 98 -digit divisor of $7^{128}+6^{128}$ obtained by elimination of "small" primes (3329 and 7329793). If we found only a few small factors we applied the SNFS; otherwise we applied the GNFS.

The numbers factored with the SNFS are listed below, and the relevant statistics appear in Table 1. Figure 1 on page 236 showed the dependence on $a$ and $b$ of the number of relations found, for the

C119 from $3^{319}-1$. Figure 2 shows the dependence on $b$ of the number of relations, and Figure 3 the average computation time per relation, for the C98 from $7^{128}+6^{128}$.

The numbers factored with the GNFS are listed below, and the relevant statistics appear in Table 2. Figure 4 shows, for the C97 from $12^{441}+1$, the effect that varying the printing bounds $W_{1}=W_{2}$ has on the time needed for the square root step. The square root program needs the full factorization of both integers $F_{i}(a, b)$ of the relations in $\mathcal{S}$. Primes larger than $W_{i}$ are printed in the input file, while smaller ones have to be found with trial division. With large printing bounds trial divisions become very time-consuming. This dependence explains the varying results for square root timings in the preceding tables.

C98 from $7^{128}+6^{128}$, factored into two primes of 49 and 50 digits, 1066005182572362964065225039667736303849504290049 and 57198455483606292671608097698721205785907158143489.

C106 from $2^{543}-1$, factored into two primes of 42 and 64 digits, 534955385319592511227419175872576025063351 and 2307880312514050317434773233753379487634082230810808744501836223.

C119 from $3^{319}-1$, factored into two primes of 41 and 79 digits,
26425387421490471188793734763177943613329 and
2742852582186302944625818325876362221386271693831170052360109818510078843584437.

Note that $g_{1}(x)=x^{10}+x^{9}+\cdots+x+1, g_{2}(x)=x-3^{29}$, and $m=3^{29}$ satisfy the requirements. Now express $g_{1}(x) / x^{5}$ as a polynomial in $x+(1 / x)$.
C123 from $2^{511}-1$, factored into two primes of 57 and 67 digits,
144780974187086260903935034761413745643636578290924150417 and
2537599745025519134156761164267591913521835535529224725592538658153.

Note that $g_{1}(x)=x^{6}+x^{5}+\cdots+x+1, g_{2}(x)=x-2^{73}$, and $m_{g}=2^{73}$ satisfy the requirements. Expressing $g_{1}(x) / x^{3}$ as a polynomial in $x+(1 / x)$ yields $h_{1}(x)=x^{3}+x^{2}-2 x-1$ with $m_{h}=2^{73}+2^{-73}$. Use
$m_{h}=2\left(\left(2^{36}+2^{-37}\right)^{2}-1\right)$ to rewrite $h_{1}(x)$ as a polynomial in $2^{36}+2^{-37}$. This yields
$k_{1}(x)=8 x^{6}-20 x^{4}+12 x^{2}-1$ and $m_{k}=2^{36}+2^{-37}$. Finally $f_{1}(x)=8 k_{1}(x / 2)$, so $m_{f}=2^{37}+2^{-36}$.
$\mathbf{C 1 3 5}$ from $73^{73}+1$, factored into two primes of 55 and 80 digits,
4596369165585291112352829637852339157090144708807832677 and
30968642349372168556028560872092954382289515106526406243465921744900645899326733.

The C162 $\left(12^{151}-1\right) / 11$, factored into two primes of 44 and 119 digits,
16537237851564688924261407041648853990657743 and $49717867800323378818763399005960016 \backslash$
487476598349539211569747005759153228241911167043200927016884285731030248831349126419.

Numbers factored with the SNFS.

|  | C 98 | C 106 | C 119 |
| :--- | :---: | :---: | :---: |
| factor of | $7^{128}+6^{128}$ | $2^{543}-1$ | $3^{319}-1$ |
| $f_{1}(x)$ | $x^{4}+1$ | $4 x^{4}+2 x^{2}+1$ | $x^{5}+x^{4}-4 x^{3}-3 x^{2}+3 x+1$ |
| $f_{2}(x)$ | $6^{32} x-7^{32}$ | $x-2^{90}$ | $3^{29} x-3^{58}-1$ |
| $m$ | $7^{32} 6^{-32} \mathrm{mod} n$ | $2^{90}$ | $3^{29}+3^{-29} \mathrm{mod} n$ |
| sieving region | $\|a\| \leq 2 \cdot 10^{6}$ | $\|a\| \leq 3.5 \cdot 10^{5}$ | $\|a\| \leq 16 \cdot 10^{5}, 1 \leq b \leq 103000 ;$ |
|  | $1 \leq b \leq 16 \cdot 10^{3}$ | $1 \leq b \leq 10^{5}$ | $\|a\| \leq 12 \cdot 10^{5}, 103001 \leq b \leq 345000$ |
| $B_{1}$ | $1.6 \cdot 10^{6}$ | $5 \cdot 10^{5}$ | $10^{6}$ |
| $B_{2}$ | $1.6 \cdot 10^{6}$ | $8.1 \cdot 10^{5}$ | $1.4 \cdot 10^{6}$ |
| $L_{1}$ | $3 \cdot 10^{7}$ | $12 \cdot 10^{6}$ | $2 \cdot 10^{7}$ |
| $L_{2}$ | $3 \cdot 10^{7}$ | $12 \cdot 10^{6}$ | $2.5 \cdot 10^{7}$ |
| sieving time | 450 hours | 250 hours | 800 hours |
| \# sieving rel. | $2,337,618$ | $1,106,949$ | $2,221,686$ |
| \# filter rel. $/$ sets | $982,672 / 587,076$ | $264,583 / 126,254$ | $774,265 / 349,961$ |
| matrix size | $539,020 \times 620,650$ | $128,546 \times 133,738$ | $348,852 \times 367,182$ |
| bl. Lanczos time | 74 min | 6 min | 38 min |
| square root time | 69 min | 47 min | 62 min |
| \# trials | 1 | 2 | 1 |


|  | C 123 | C 135 | C 162 |
| :--- | :---: | :---: | :---: |
| factor of | $2^{511}-1$ | $73^{73}+1$ | $12^{151}-1$ |
| $f_{1}(x)$ | $x^{6}-10 x^{4}+24 x^{2}-8$ | $x^{5}+73^{2}$ | $12 x^{5}-1$ |
| $f_{2}(x)$ | $2^{36} x-2^{73}-1$ | $x-73^{15}$ | $x-12^{30}$ |
| $m$ | $2^{37}+2^{-36} \bmod n$ | $73^{15}$ | $12^{30}$ |
| sieving region | $\|a\| \leq 6 \cdot 10^{5}$ | $\|a\| \leq 2 \cdot 10^{6}$ | $?$ |
| $B_{1}$ | $1 \leq b \leq 37 \cdot 10^{4}$ | $1 \leq b \leq 2.6 \cdot 10^{5}$ |  |
| $B_{2}$ | $15 \cdot 10^{5}$ | $2 \cdot 10^{6}$ | $?$ |
| $L_{1}$ | $11 \cdot 10^{5}$ | $2 \cdot 10^{6}$ | $?$ |
| $L_{2}$ | $3 \cdot 10^{7}$ | $3 \cdot 10^{7}$ | $10^{8}$ |
| sieving time | $2.5 \cdot 10^{7}$ | $3 \cdot 10^{7}$ | $10^{8}$ |
| \# sieving rel. | 700 hours | 2150 hours | see caption |
| \# filter rel. / sets | $420,896 / 222,014$ | $1,154,111 / 583,631$ | $8.98 \cdot 10^{6}$ |
| matrix size | $430,018 \times 439,058$ | $581,870 \times 590,573$ | $1,807,808 / 822,361$ |
| bl. Lanczos time | 97 hours | 92 min | $828,077 \times 833,017$ |
| square root time | 13 hours | 130 min | 205 min |
| \# trials | 1 | 1 | 10.5 hours |

TABLE 1. Statistics of SNFS runs. The square root timings are given for one dependency. "Sieving time" is a rough estimate on an SGI Indy workstation ( 100 MHz R4000SC processor), except for the C162, where it represents 8 weeks of idle time on 30 workstations at Oregon State University, Corwallis (USA). "Block Lanczos time" is time on one processor of a Cray C98, except for the C123, where it is on one processor of an SGI Challenge ( 150 MHz R4400SC processor). "Square root time" is always on one 150 MHz R 4400 SC processor of an SGI Challenge. A question mark indicates that records have been lost. See also Figure 1 for the C119 and Figures 2 and 3 for the C98.

## ACKNOWLEDGEMENTS

I am particularly grateful to P. L. Montgomery for helping me understand the Number Field Sieve and for sharing with me his NFS program, developed partially by A. K. Lenstra and Oregon State University. I thank H. W. Lenstra, Jr., P. L. Montgomery, H. J. J. te Riele, and R. Tijdeman for reading the paper and for suggesting several im-
provements. I also thank the referee for remarks that helped improve the presentation in Section 10.

This work was sponsored by the Stichting Nationale Computerfaciliteiten (National Computing Facilities Foundation, NCF) for the use of supercomputer facilities, with financial support from the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (Netherlands Organization for Scientific Research, NWO).
$\mathbf{C 8 7}$ from $72^{99}+1$, factored into two primes of 28 and 59 digits,
8097540789168990910686588841 and 16798259634305265460776431824065479289925925226507262384081.
$151128965885928672997 x^{2}-1445482064037103573030 x-8289110711415033862728$
$1019302053499425445447 x^{2}+6959100565519575386428 x-73926156784894075263915$
129714918197797463455765215830559869325936651203090005723712501519760440714275563588906
$729879993 x^{2}+195749170187826932902 x-7755788285642477233065859266474816$
$2173217947 x^{2}-7907605467511404257356 x-23092862466613462049239427852471175$
91584635604178938356370529300221754097732874797358880201218639497144310994413684200466
C97 from $12^{441}-1$, factored into two primes of 44 and 53 digits, 80563595863991546010661115881806884967004027 and 34988882229141372324919456579820527956234275361821641.
$-30177438253 x^{2}+435491824510787395448959 x+3562987893564232236048460618728361350$
$-536199277805 x^{2}-4803940130457949745361008 x+63307942819437806775380979481980102858$
287932991427719831471392473244890825244326896138836873469449401613203513827866583225758151286534
C105 from $3^{367}-1$, factored into two primes of 52 and 54 digits,
1511495257840070716998865694022937935039928231350493 and
501953997389244528404247906279090654105468962124251929.
$342910527737 x^{2}+86817069333519465483641612 x+540759062604782971357139536186424874771$
$1242060255079 x^{2}-913049273181768816962553218 x+129128767300065233631168229536267982420800$
229143590555869469062115013538557681923164235754266217765793563500275674926893987223245481401160544005942
C106 from $12^{157}+1$ was factored into two primes of 43 and 63 digits,
5907648479166682692493875790809420049637197 and
780034754594869441475241918877754732415095147758932024091338271.
$190030476113 x^{2}-101643163734436736066960294 x-32430287560495976143910317159823376255144$
$-785083260639 x^{2}-401968646051742270344280172 x+164086080001456034179238766543256687713827$
1790044128757262576848153412133765937899097888814377815816769105476827696665209945565825606429787588581699
C107 from $6^{223}+1$, factored into 2 primes of 48 and 59 digits,
835790655259197870586319955878764574136853697743 and
16660359818385555664068601979347465063709118473978852096987.
$-540161776283 x^{2}-425704283028714253779269315 x-46786964108579179806101863478910720071558$
$-241799514805 x^{2}+763119703166287854853198889 x-311653994359418670319775330136434513506986$
12637530599467776761853128412624277137347729851839924048392287605249253270797264409813230653725405155484892

Numbers factored with the GNFS. After each factorization we give the polynomials $f_{1}$ and $f_{2}$ used in the run, and the integer $m$. The C87 was factored twice, once with classical sieving (first set of values of $f_{1}, f_{2}, m$ ) and once with lattice sieving. See also Figure 4 for the C97.

|  | C 87 | C 87 | C 97 |
| :--- | :---: | :---: | :---: |
| factor of | $72^{99}+1$ | $72^{99}+1$ | $12^{441}+1$ |
| sieving method | classical | lattice | lattice |
| sieving region | $\|a\| \leq 2 \cdot 10^{6}, 1 \leq b \leq 48 \cdot 10^{4}$ | $\|a\| \leq 7.5 \cdot 10^{12}, b=1$ | $\|a\| \leq 25 \cdot 10^{12}, b=1$ |
| $B_{1}=B_{2}$ | $1.6 \cdot 10^{6}$ | $10^{6}$ | $2.2 \cdot 10^{6}$ |
| $L^{(l)}$ | $\left(L_{1}=4 \cdot 10^{7}\right)$ | $10^{6}$ | $10 \cdot 10^{6}$ |
| $L^{(u)}$ | $\left(L_{2}=4 \cdot 10^{7}\right)$ | $2.346 \cdot 10^{6}$ | $L^{(u)}=24 \cdot 10^{6}$ |
| sieving time | 2100 hours | 1500 hours | 3500 hours |
| \# sieving rel. | $3,480,325$ | 521,901 | $3,599,014$ |
| \# filter rel. / sets | $741,930 / 338,580$ | $426,241 / 409,699$ | $1,247,094 / 604,205$ |
| matrix size | $364,215 \times 366,907$ | $273,475 \times 437,441$ | $637,711 \times 644,950$ |
| bl. Lanczos time | $41 \min ^{(2)}$ | $26 \min ^{(2)}$ | $123 \min ^{(2)}$ |
| square root time | $128 \min ^{(3)}$ | $36 \min ^{(3)}$ | $78 \min ^{(3)}$ |
| \# trials | 1 | 2 | 2 |


|  | C 105 | C 106 | C 107 |
| :--- | :---: | :---: | :---: |
| factor of | $3^{367}-1$ | $12^{157}+1$ | $6^{223}+1$ |
| sieving method | lattice | lattice | lattice |
| sieving region | $\|a\| \leq 7.5 \cdot 10^{14}, b=1$ | $\|a\| \leq 10^{15}, b=1$ | $\|a\| \leq 10^{15}, b=1$ |
| $B_{1}=B_{2}$ | $1.6 \cdot 10^{6}$ | $2.7 \cdot 10^{6}$ | $2.9 \cdot 10^{6}$ |
| $L^{(l)}$ | $23 \cdot 10^{6}$ | $2.7 \cdot 10^{7}$ | $2.72 \cdot 10^{7}$ |
| $L^{(u)}$ | $30 \cdot 10^{6}$ | $3 \cdot 10^{7}$ | $3 \cdot 10^{7}$ |
| sieving time | see caption | 11900 hours | 11200 hours |
| \# sieving rel. | $3.59 \cdot 10^{6}$ | $3,272,224$ | $3,098,987$ |
| \# filter rel. / sets | $? / 1,218,633$ | $2,151,431 / 1,191,636$ | $2,152,685 / 1,155,270$ |
| matrix size | $1,284,719 \times 1,294,861$ | $1,266,098 \times 1,295,043$ | $1,226,577 \times 1,252,846$ |
| bl. Lanczos time | 439 min $^{(2)}$ | 423 min $^{(2)}$ | $421 \mathrm{~min}^{(2)}$ |
| square root time | 4.8 hours $^{(3)}$ | 2.0 hours $^{(4)}$ | $2.1 \mathrm{hours}^{(3)}$ |
| \# trials | 1 | 1 | 5 |

TABLE 2. Statistics of GNFS runs. The square root timings are given for one dependency. "Sieving time" is a rough estimate on an SGI Indy workstation ( 100 MHz R4000SC processor), except for the C105, where it represents 8 weeks of idle time on 40 workstations at Oregon State University, Corwallis (USA). "Block Lanczos time" is time on one processor of a Cray C98. "Square root time" is on one processor of an SGI Challenge ( 150 MHz R4400SC processor), except for the C106, where the clock rate was 200 MHz . A question mark indicates that records have been lost. See also Figure 4 for the C97.

## REFERENCES

[Adleman 1991] L. M. Adleman, "Factoring numbers using singular integers", pp. 64-71 in Proceedings 23 rd Annual ACM Symposium on Theory of Computing (New Orleans, 1991), ACM, New York, 1991.
[Batut et al. 1995] C. Batut, D. Bernardi, H. Cohen
and M. Olivier, User's Guide to Pari-GP. This manual is part of the program distribution, available by anonymous ftp from the host megrez.math. u-bordeaux.fr.
[Buhler et al. 1993] J. P. Buhler, H. W. Lenstra, Jr., and C. Pomerance, "Factoring integers with the number field sieve", pp. 50-94 in [Lenstra et al. 1993a].


FIGURE 2. Number of relations found in each range of $500 b$-values for the factorization of the C98 from $7^{128}+6^{128}$.

FIGURE 4. Square root timings (on one 150 MHz R4400SC processor of an SGI Challenge) for the C97 from $12^{441}+1$.


FIGURE 3. Approximate average time per relation needed in the factorization run of the C98 from $7^{128}+6^{128}$, again as a function of $b$. Times are on an SGI Indy workstation with a 100 MHz R 4000 SC processor.

FIGURE 5. All matrices from the experiments had close to the average of 30.3 nonzero elements per row. Here we show the dependence of the block Lanczos timings on the product of the number of rows and columns.
[Buhler et al. 1994] J. P. Buhler, P. L. Montgomery, R. Robson, and R. Ruby, "Technical report implementing the number field sieve", Oregon State University, Corwallis, OR, 1994.
[Cohen 1993] H. Cohen, A Course in Computational Algebraic Number Theory, Springer, Berlin, 1993.
[Coppersmith 1993] D. Coppersmith, "Solving linear equations over GF(2): Block Lanczos algorithm", Linear Algebra and its Applications 192 (1993), 3360.
[Couveignes 1993] J.-M. Couveignes, "Computing a square root for the number field sieve", pp. 95-102 in [Lenstra et al. 1993a].
[Frobenius 1896] F. G. Frobenius. "Über Beziehungen zwischen den Primidealen eines algebraischen Körpers und den Substitutionen seiner Gruppe", Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften zu Berlin (1896), pp. 689-703. Reprinted in Gesammelte Abhandlungen, Band II, Springer, Berlin, 1968.
[Golliver et al. 1994] R. A. Golliver, A. K. Lenstra, and K. S. McCurley, "Lattice sieving and trial division", pp. 18-27 in Algorithmic Number Theory (edited by L. M. Adleman and M.-D. Huang), Lecture Notes in Comp. Sci. 877, Springer, Berlin, 1994.
[Knuth 1981] D. E. Knuth, The Art of Computer Programming, vol. 2: Seminumerical Algorithms, 2nd ed., Addison-Wesley, Reading, MA, 1981 (or 3rd ed., 1997).
[LaMacchia et al. 1991] B. A. LaMacchia and A. M. Odlyzko, "Solving large sparse linear systems over finite fields", pp. 109-133 Advances in Cryptology: CRYPTO ' 90 (edited by A. J. Menezes and S. A. Vanstone), Lecture Notes in Comp. Sci. 537, Springer, Berlin, 1991.
[Lang 1970] S. Lang, Algebraic Number Theory, Addison-Wesley, Reading, MA, 1970.
[Lenstra et al. 1993a] A. K. Lenstra and H. W. Lenstra, Jr., The Development of the Number Field Sieve, Lecture Notes in Math. 1554, Springer, Berlin, 1993.
[Lenstra et al. 1982] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász, "Factoring polynomials with rational coefficients", Mathematische Annalen, 261 (1982), 515-534.
[Lenstra et al. 1993b] A. K. Lenstra, H. W. Lenstra, Jr., M. S. Manasse, and J. M. Pollard, "The number field sieve", pp. 11-42 in [Lenstra et al. 1993a].
[Lenstra et al. 1993c] A. K. Lenstra, H. W. Lenstra, Jr., M. S. Manasse, and J. M. Pollard, "The factorization of the ninth Fermat number", Mathematics of Computation, 61 (1993), 319-349.
[Montgomery 1994] Peter L. Montgomery, "Square roots of products of algebraic numbers", pp. 567571 in Mathematics of Computation 1943-1993: A Half-Century of Computational Mathematics (edited by Walter Gautschi), Proceedings of Symposia in Applied Mathematics, American Math. Soc., Providence, 1994. Long version to appear.
[Montgomery 1995] Peter L. Montgomery, "A block Lanczos algorithm for finding dependencies over GF(2)", pp. 106-120 in Advances in Cryptology: Eurocrypt '95 (edited by L. C. Guillou and J.J. Quisquater), Lecture Notes in Comp. Sci. 921, Springer, Berlin, 1995.
[Neukirch 1992] J. Neukirch, Algebraische Zahlentheorie, Springer, Berlin, 1992.
[Pollard 1993a] J. M. Pollard, "Factoring with cubic integers", pp. 4-10 in [Lenstra et al. 1993a].
[Pollard 1993b] J. M. Pollard, "The lattice sieve", pp. 43-49 in [Lenstra et al. 1993a].
[Pomerance 1985] C. Pomerance, "The quadratic sieve factoring algorithm", pp. 169-182 in Advances in Cryptology: Eurocrypt' 84 (edited by T. Beth, N. Cot, and I. Ingemarsson), Lecture Notes in Comp. Sci. 209, Springer, New York, 1985.
[Riesel 1985] H. Riesel, Prime Numbers and Computer Methods for Factorization, Birkhäuser, Boston, 1985.
[Standish 1980] T. A. Standish, Data Structure Techniques, Addison-Wesley, Reading, MA, 1980.

Marije Elkenbracht-Huizing, CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands (marije@cwi.nl)
Received August 24, 1995; accepted in revised form April 9, 1996

