# Improved Stage 2 to $\mathrm{P} \pm 1$ Factoring Algorithms 

Peter L. Montgomery ${ }^{1}$ and Alexander Kruppa ${ }^{2}$<br>${ }^{1}$ Microsoft Research, One Microsoft Way, Redmond, WA 98052 USA.<br>pmontgom@cwi.nl<br>${ }^{2}$ LORIA, Campus Scientifique, BP 239, 54506 Vandœuvre-lès-Nancy Cedex, France.<br>kruppaal@loria.fr


#### Abstract

Some implementations of stage 2 of the $\mathrm{P}-1$ method of factorization use convolutions. We describe a space-efficient implementation, allowing convolution lengths around $2^{23}$ and stage 2 limit around $10^{16}$ while attempting to factor 230 -digit numbers on modern PC's. We use the discrete cosine transform to multiply reciprocal polynomials. We present adjustments for the $\mathrm{P}+1$ algorithm. We list some new findings. Key words. Integer factorization, convolution, discrete cosine transform, DCT-II, discrete Fourier transform, $\mathrm{P}+1, \mathrm{P}-1$, multipoint polynomial evaluation, reciprocal polynomials.


## 1 Introduction

John Pollard introduced the $\mathrm{P}-1$ algorithm for factoring an odd composite integer $N$ in 1974 [?, section 4]. It hopes that some prime factor $p$ of $N$ has smooth $p-1$. It picks $b_{0}$ coprime to $N$ and outputs $b_{1}=b_{0}^{e} \bmod N$ for some positive exponent $e$. This exponent might be divisible by all prime powers below a bound $B_{1}$. Stage 1 succeeds if $(p-1) \mid e$, in which case $b_{1} \equiv 1(\bmod p)$ by Fermat's little theorem. The algorithm recovers $p$ by computing $\operatorname{gcd}\left(b_{1}-1, N\right)$ (except in rare cases where this GCD is composite). When this GCD is 1 , we hope that $p-1=q n$ where $n$ divides $e$ and $q$ is not too large. Then

$$
\begin{equation*}
b_{1}^{q} \equiv\left(b_{0}^{e}\right)^{q}=b_{0}^{e q}=\left(b_{0}^{n q}\right)^{e / n}=\left(b_{0}^{p-1}\right)^{e / n} \equiv 1^{e / n}=1 \quad(\bmod p) \tag{1}
\end{equation*}
$$

so $p$ divides $\operatorname{gcd}\left(b_{1}^{q}-1, N\right)$. Stage 2 of $\mathrm{P}-1$ tries to find $p$ when $q>1$ but $q$ is not too large. The search bound for $q$ is called $B_{2}$.

Pollard [?] tries each prime $q$ in [ $B_{1}, B_{2}$ ] individually. If $q_{1}$ and $q_{2}$ are successive primes, then look up $b_{1}^{q_{2}-q_{1}} \bmod N$ in a small table. Given $b_{1}^{q_{1}} \bmod N$, form $b_{1}^{q_{2}} \bmod N$ and test $\operatorname{gcd}\left(b_{1}^{q_{2}}-1, N\right)$. He observes that one can combine the GCD tests: if either $p \mid \operatorname{gcd}(x, N)$ or $p \mid \operatorname{gcd}(y, N)$, then $p \mid \operatorname{gcd}(x y \bmod N, N)$. His stage 2 cost is two modular multiplications per $q$, one GCD with $N$ at the end, and a few multiplications to build the table.

Montgomery [?] uses two sets $S_{1}$ and $S_{2}$, such that each prime $q$ in [ $B_{1}, B_{2}$ ] divides a nonzero difference $s_{1}-s_{2}$ where $s_{1} \in S_{1}$ and $s_{2} \in S_{2}$. He forms $b_{1}^{s_{1}}-b_{1}^{s_{2}}$ using two table look-ups, saving one modular multiplication per $q$. Sometimes one
$s_{1}-s_{2}$ works for multiple $q$. Montgomery adapts his scheme to Hugh Williams's P +1 method and Hendrik Lenstra's elliptic curve method (ECM).

These changes lower the constant of proportionality, but stage 2 still uses $O\left(\pi\left(B_{2}\right)-\pi\left(B_{1}\right)\right)$ (number of primes between $B_{1}$ and $B_{2}$ ) operations modulo $N$.

The end of [?] suggests an FFT continuation to $\mathrm{P}-1$. Silverman [?, p. 844] implements it, using a circular convolution to evaluate a polynomial along a geometric progression. Montgomery's dissertation [?] describes an FFT continuation to ECM. He takes the GCD of two polynomials. Zimmermann [?] implements another FFT continuation to ECM, based on evaluating a polynomial at arbitrary points. Zimmermann adapts his implementation to $P \pm 1$ methods.

Like [?], we evaluate a polynomial along geometric progressions. We exploit patterns in its roots to generate its coefficients quickly. We aim for low memory overhead, saving it for convolution inputs and outputs (which are elements of $\mathbb{Z} / N \mathbb{Z})$. Using memory efficiently lets us raise the convolution length $\ell$. Many intermediate results are reciprocal polynomials, which need about half the storage and can be multiplied using the discrete cosine transform.

Doubling $\ell$ costs slightly over twice as much time per convolution, but each longer convolution extends the search for $q$ (and effective $B_{2}$ ) fourfold. Silverman's 1989 implementation used 42 megabytes and allowed 250-digit inputs. It repeatedly evaluated a polynomial of degree 15360 at $8 \cdot 17408$ points in geometric progression, using $\ell=32768$. This enabled him to achieve $B_{2} \approx 10^{10}$.

Today's (2007) PC memories are 100 times as large as that used in [?]. With this extra memory, we achieve $\ell=2^{23}$, a growth factor of 256 . With the same number of convolutions (individually longer lengths but running on faster hardware) our $B_{2}$ advances by a factor of $256^{2} \approx 6.6 e 4$.

Section ?? gives some new results, including a 60 -digit $\mathrm{P}+1$ factor.

## $2 \mathrm{P}+1$ Algorithm

Hugh Williams [?] introduced a $\mathrm{P}+1$ factoring algorithm. It finds a prime factor $p$ of $N$ when $p+1$ (rather than $p-1$ ) is smooth. It is modeled after $\mathrm{P}-1$.

One variant of the $\mathrm{P}+1$ algorithm chooses $P_{0} \in \mathbb{Z} / N \mathbb{Z}$ and lets the indeterminate $\alpha_{0}$ be a zero of the quadratic $\alpha_{0}^{2}-P_{0} \alpha_{0}+1$. We hope this quadratic is irreducible modulo $p$. If so, its second root in $\mathbb{F}_{p^{2}}$ will be $\alpha_{0}^{p}$. The product of its roots is the constant term 1 . Hence $\alpha_{0}^{p+1} \equiv 1(\bmod p)$ when we choose well.

Stage 1 of the $\mathrm{P}+1$ algorithm computes $P_{1}=\alpha_{1}+\alpha_{1}^{-1}$ where $\alpha_{1} \equiv \alpha_{0}^{e}$ $(\bmod N)$ for some exponent $e$. If $\operatorname{gcd}\left(P_{1}^{2}-4, N\right)>1$, then the algorithm succeeds. Stage 2 of $\mathrm{P}+1$ hopes that $\alpha_{1}^{q} \equiv 1(\bmod p)$ for some prime $q$, not too large, and some prime $p$ dividing $N$.

Most techniques herein adapt to $\mathrm{P}+1$, but some computations take place in an extension ring, raising memory usage if we use the same convolution sizes.

### 2.1 Chebyshev Polynomials

Although the theory behind $\mathrm{P}+1$ mentions $\alpha_{0}$ and $\alpha_{1}=\alpha_{0}^{e}$, an implementation manipulates primarily values of $\alpha_{0}^{n}+\alpha_{0}^{-n}$ and $\alpha_{1}^{n}+\alpha_{1}^{-n}$ for various integers $n$ rather than the corresponding values (in an extension ring) of $\alpha_{0}^{n}$ and $\alpha_{1}^{n}$.

Let $n$ be an integer. The Chebyshev polynomials $V_{n}$ and $U_{n}$ satisfy the formal identities

$$
\begin{aligned}
V_{n}\left(X+X^{-1}\right) & =X^{n}+X^{-n} \\
\left(X-X^{-1}\right) U_{n}\left(X+X^{-1}\right) & =X^{n}-X^{-n}
\end{aligned}
$$

The use of these polynomials shortens many formulas, such as

$$
P_{1} \equiv \alpha_{1}+\alpha_{1}^{-1} \equiv \alpha_{0}^{e}+\alpha_{0}^{-e}=V_{e}\left(\alpha_{0}+\alpha_{0}^{-1}\right)=V_{e}\left(P_{0}\right) \quad(\bmod N)
$$

These polynomials have integer coefficients, so $P_{1} \equiv V_{e}\left(P_{0}\right)(\bmod N)$ is in the base ring $\mathbb{Z} / N \mathbb{Z}$ even when $\alpha_{0}$ and $\alpha_{1}$ are not.

The Chebyshev polynomials satisfy many identities, including

$$
\begin{align*}
V_{m n}(X) & =V_{m}\left(V_{n}(X)\right) \\
U_{m+n}(X) & =U_{m}(X) V_{n}(X)-U_{m-n}(X)  \tag{2}\\
U_{m+n}(X) & =V_{m}(X) U_{n}(X)+U_{m-n}(X) \\
V_{m+n}(X) & =V_{m}(X) V_{n}(X)-V_{m-n}(X)  \tag{3}\\
V_{m+n}(X) & =\left(X^{2}-4\right) U_{m}(X) U_{n}(X)+V_{m-n}(X)
\end{align*}
$$

## 3 Overview of Stage 2 Algorithm

Our algorithm performs multipoint evaluation of polynomials by convolutions. Its inputs are the output of stage $1\left(b_{1}\right.$ for $\mathrm{P}-1$ or $P_{1}$ for $\left.\mathrm{P}+1\right)$, and the desired stage 2 interval $\left[B_{1}, B_{2}\right]$.

It chooses an odd integer $P$ with large $P / \phi(P)$, a convolution length $\ell_{\max }$ and a factorization $\phi(P)=s_{1} s_{2}$ so that $s_{1}$ is even and close to $\ell_{\max } / 2$, perhaps $0.3 \leq s_{1} / \ell_{\max } \leq 0.7$. We require $\ell_{\max }>s_{1}$. It will do $s_{2}$ convolutions of length $\ell_{\text {max }}$; each evaluates a polynomial of degree $s_{1}$ on $\ell_{\max }-s_{1}$ points.

Using a factorization of $(\mathbb{Z} / P \mathbb{Z})^{*}$ as described in section ??, it constructs two sets $S_{1}$ and $S_{2}$ of integers such that
(a) $\left|S_{1}\right|=s_{1}$ and $\left|S_{2}\right|=s_{2}$.
(b) $S_{1}$ is symmetric around 0 . If $k \in S_{1}$, then $-k \in S_{1}$.
(c) If $k \in \mathbb{Z}$ and $\operatorname{gcd}(k, P)=1$, then there exist unique $k_{1} \in S_{1}$ and $k_{2} \in S_{2}$ such that $k \equiv k_{1}+k_{2}(\bmod P)$.

Once $S_{1}$ and $S_{2}$ are chosen, it computes the coefficients of

$$
\begin{equation*}
f(X)=X^{-s_{1} / 2} \prod_{k_{1} \in S_{1}}\left(X-b_{1}^{2 k_{1}}\right) \bmod N \tag{4}
\end{equation*}
$$

by the method in section ??. Since $S_{1}$ is symmetric around zero, this $f(X)$ is symmetric in $X$ and $1 / X$.

For each $k_{2} \in S_{2}$ it evaluates (the numerators of) all

$$
\begin{equation*}
f\left(b_{1}^{2 k_{2}+(2 m+1) P}\right) \bmod N \tag{5}
\end{equation*}
$$

for $\ell_{\max }-s_{1}$ consecutive values of $m$ as described in section ??, and checks the product of these outputs for a nontrivial GCD with $N$.

For the $\mathrm{P}+1$ method, replace (??) by $f(X)=X^{-s_{1} / 2} \prod_{k_{1} \in S_{1}}\left(X-\alpha_{1}^{2 k_{1}}\right) \bmod$ $N$. Similarly, replace $b_{1}$ by $\alpha_{1}$ in (??). The polynomial $f$ is still over $\mathbb{Z} / N \mathbb{Z}$, but the multipoint evaluation works in an extension ring. See section ??.

## 4 Justification

Let $p$ be an unknown prime factor of $N$. As in (??), assume $b_{1}^{q} \equiv 1(\bmod p)$ where $q$ is not too large, and $\operatorname{gcd}(q, 2 P)=1$.

The selection of $S_{1}$ and $S_{2}$ ensures there exist $k_{1} \in S_{1}$ and $k_{2} \in S_{2}$ such that $(q-P) / 2 \equiv k_{1}+k_{2}(\bmod P)$. That is,

$$
\begin{equation*}
q=P+2 k_{1}+2 k_{2}+2 m P=2 k_{1}+2 k_{2}+(2 m+1) P \tag{6}
\end{equation*}
$$

for some integer $m$. We can bound $m$ knowing bounds on $q, k_{1}, k_{2}$, detailed in the next section. Both $b_{1}^{ \pm 2 k_{1}}$ are roots of $f(\bmod p)$. Hence

$$
\begin{equation*}
f\left(b_{1}^{2 k_{2}+(2 m+1) P}\right)=f\left(b_{1}^{q-2 k_{1}}\right) \equiv f\left(b_{1}^{-2 k_{1}}\right) \equiv 0 \quad(\bmod p) \tag{7}
\end{equation*}
$$

For the $\mathrm{P}+1$ method, if $\alpha_{1}^{q} \equiv 1(\bmod p)$, then $(? ?)$ evaluates $f$ at $X=$ $\alpha_{1}^{2 k_{2}+(2 m+1) P}=\alpha_{1}^{q-2 k_{1}}$. The factor $X-\alpha_{1}^{-2 k_{1}}$ of $f(X)$ evaluates to $r^{-2 k_{1}}\left(\alpha_{1}^{q}-1\right)$, which is zero modulo $p$ even in the extension ring.

## 5 Selection of $S_{1}$ and $S_{2}$

Let "+" of two sets denote the set of sums. By the Chinese Remainder Theorem,

$$
\begin{equation*}
(\mathbb{Z} /(m n) \mathbb{Z})^{*}=n(\mathbb{Z} / m \mathbb{Z})^{*}+m(\mathbb{Z} / n \mathbb{Z})^{*} \text { if } \operatorname{gcd}(m, n)=1 \tag{8}
\end{equation*}
$$

This is independent of the representatives: if $S \equiv(\mathbb{Z} / m \mathbb{Z})^{*}(\bmod m)$ and $T \equiv$ $(\mathbb{Z} / n \mathbb{Z})^{*}(\bmod n)$, then $n S+m T \equiv(\mathbb{Z} /(m n) \mathbb{Z})^{*}(\bmod m n)$. For prime powers, $\left(\mathbb{Z} / p^{k} \mathbb{Z}\right)^{*}=(\mathbb{Z} / p \mathbb{Z})^{*}+\sum_{i=1}^{k-1} p^{i}(\mathbb{Z} / p \mathbb{Z})$.

We choose $S_{1}$ and $S_{2}$ so that $S_{1}+S_{2} \equiv(\mathbb{Z} / P \mathbb{Z})^{*}(\bmod P)$ which ensures that all values coprime to $P$, in particular all primes, in the stage 2 interval are covered. One way uses a factorization $m n=P$ and (??). Other choices are available by factoring individual $(\mathbb{Z} / p \mathbb{Z})^{*}, p \mid P$, into smaller sets of sums.

Let $R_{n}=\{2 i-n-1: 1 \leq i \leq n\}$ be an arithmetic progression centered at 0 of length $n$ and common difference 2 . For odd primes $p$, a set of representatives of $(\mathbb{Z} / p \mathbb{Z})^{*}$ is $R_{p-1}$. Its cardinality is composite for $p \neq 3$ and the set can be
factored into arithmetic progressions of prime length by $R_{n m}=R_{n}+n R_{m}$. If $p \equiv 3(\bmod 4)$, alternatively $\frac{p+1}{4} R_{2}+\frac{1}{2} R_{(p-1) / 2}$ can be chosen as a set of representatives with smaller absolute values.

When evaluating (??) for all $m_{1} \leq m<m_{2}$ and $k_{2} \in S_{2}$, the highest exponent coprime to $P$ that is not covered at the low end of the stage 2 range will be $2 \max \left(S_{1}+S_{2}\right)+\left(2 m_{1}-1\right) P$. Similarly, the smallest value at the high end of the stage 2 range not covered is $2 \min \left(S_{1}+S_{2}\right)+\left(2 m_{2}+1\right) P$. Hence, for a given choice of $P, S_{1}, S_{2}, m_{1}$ and $m_{2}$, all primes in $\left[\left(2 m_{1}-1\right) P+2 \max \left(S_{1}+S_{2}\right)+1\right.$, $\left.\left(2 m_{2}+1\right) P+2 \min \left(S_{1}+S_{2}\right)-1\right]$ are covered.

For choosing a value of $P$ which covers a desired $\left[B_{1}, B_{2}\right.$ ] interval, we can test candidate $P$ from a table. This table could contain values so that $P$ and $\phi(P)$ are increasing, and each $P$ is maximal for its $\phi(P)$. We can select those $P$ which, in order, cover the desired $\left[B_{1}, B_{2}\right]$ interval with the user-specified $\ell_{\max }$, minimize $s_{2} \cdot \ell$ and maximize $\left(2 m_{2}+1\right) P+2 \min \left(S_{1}+S_{2}\right)$.

For example, to cover the interval $[1000,500000]$ with $\ell_{\max }=512$, we might choose $P=1155, s_{1}=240, s_{2}=2, m_{1}=-1, m_{2}=271$. With $S_{1}=$ $231(\{-1,1\}+\{-2,2\})+165(\{-2,2\}+\{-1,0,1\})+105(\{-3,3\}+\{-2,-1,0,1,2\})$ and $S_{2}=385\{-1,1\}$, we have $\max \left(S_{1}+S_{2}\right)=-\min \left(S_{1}+S_{2}\right)=2098$ and thus cover all primes in $[-3 \cdot 1155+4196+1,541 \cdot 1155-4196-1]=[732,620658]$.

## 6 Circular Convolutions and Polynomial Multiplication

Let $R$ be a commutative ring with 1 and let $\ell$ be a positive integer. A circular convolution of length $\ell$ over $R$ multiplies two polynomials $f_{1}(X)$ and $f_{2}(X)$ of degree at most $\ell-1$ in the ring $R[X]$, returning $f_{1}(X) f_{2}(X) \bmod X^{\ell}-1$.

When $\operatorname{deg}\left(f_{1}\right)+\operatorname{deg}\left(f_{2}\right)<\ell$, the convolution gives an exact product.
If $R$ has a primitive $\ell$-th root $\omega$ of unity, and if $\ell$ is not a zero divisor in $R$, then one convolution algorithm uses the Discrete Fourier Transform (DFT) [?, chapter 7]. Fix $\omega$. A forward DFT evaluates all $f_{1}\left(\omega^{i}\right)$ for $0 \leq i \leq \ell-1$. Another forward DFT evaluates all $\ell$ values of $f_{2}\left(\omega^{i}\right)$. Multiply these pointwise. Then an inverse $D F T$ interpolates to find a polynomial $f_{3} \in R[X]$ of degree at most $\ell-1$ with $f_{3}\left(\omega^{i}\right)=f_{1}\left(\omega^{i}\right) f_{2}\left(\omega^{i}\right)$ for all $i$. Return $f_{3}$.

If $\ell$ is a power of 2 and we use a Fast Fourier Transform (FFT) algorithm for the forward and inverse DFTs, then the convolution takes $O(\ell \log \ell)$ operations in a suitable ring, compared to $O\left(\ell^{2}\right)$ ring operations for the naïve algorithm.

### 6.1 Convolutions over $\mathbb{Z} / N \mathbb{Z}$

The DFT cannot be used directly when $R=\mathbb{Z} / N \mathbb{Z}$, since we don't know a suitable $\omega$. As in [?, p. 534], we consider two ways to do the convolutions.

Montgomery [?, section 4] suggests a number theoretic transform (NTT). He treats the input polynomial coefficients as integers in $[0, N-1]$ and multiplies the polynomials over $\mathbb{Z}$. The product polynomial, reduced modulo $X^{\ell}-1$, has coefficients in $\left[0, \ell(N-1)^{2}\right]$. Select some 63 -bit NTT primes $p_{j}$ such that $\prod_{j} p_{j}>\ell(N-1)^{2}$. Require each $p_{j} \equiv 1(\bmod \ell)$, so a primitive $\ell$-th root of unity
$\omega_{j}$ modulo $p_{j}$ exists. Do the convolution modulo each $p_{j}$ and use the Chinese Remainder Theorem (CRT) to determine the product over $\mathbb{Z}$ modulo $X^{\ell}-1$. Reduce this product modulo $N$. Montgomery's dissertation [?, chapter 8 ] describes these computations in detail.

The convolution code should supply functions to (1) zero a DFT buffer (or a DCT buffer - see section ??), (2) add one residue modulo $N$ to a DFT buffer at a time, i.e. reduce it modulo the NTT primes, (3) perform a forward DFT on a buffer, (4) compute a point-wise product of the DFT transforms in two buffers, overwriting an input, and performing the inverse DFT on the product, and (5) extracting residues modulo $N$ from the convolution product by performing the CRT computation and reduction modulo $N$.

Another method for computing convolutions uses fast integer multiplication. See section??.

### 6.2 Reciprocal Laurent Polynomials and Discrete Cosine Transform

Define a reciprocal Laurent polynomial (RLP) in $X$ to be an expansion $a_{0}+$ $\sum_{j=1}^{d} a_{j} V_{j}\left(X+X^{-1}\right)=a_{0}+\sum_{j=1}^{d} a_{j} \cdot\left(X^{j}+X^{-j}\right)$ for scalars $a_{j}$ in a ring. It is monic if $a_{d}=1$. It is said to have degree $2 d$ if $a_{d} \neq 0$. The degree is always even.

A monic RLP of degree $2 d$ fits in $d$ coefficients (excluding the leading 1 ).
While manipulating RLPs of degree at most $2 d$, the standard basis is $\{1\} \cup$ $\left\{X^{j}+X^{-j}: 1 \leq j \leq d\right\}=\{1\} \cup\left\{V_{j}(Y): 1 \leq j \leq d\right\}$ where $Y=X+X^{-1}$.

Let $\ell$ be a positive integer. Let $R$ be a ring in which $2 \ell$ is not a zero divisor Assume $R$ has a primitive $4 \ell$-th root $\omega_{\ell}$ of unity for which $\omega_{\ell}^{2 \ell}=-1$. Set $\gamma_{\ell}=$ $\omega_{\ell}+\omega_{\ell}^{-1}$. Then $V_{\ell}\left(\gamma_{\ell}\right)=0$. We omit the subscript $\ell$ when it is clear from context.

Let $Q(X)=q_{0}+\sum_{j=1}^{\ell-1} q_{j}\left(X^{j}+X^{-j}\right)=q_{0}+\sum_{j=1}^{\ell-1} q_{j} V_{j}(X+1 / X)$ be a RLP of degree at most $2 \ell-2$. Define its discrete cosine transform of length $\ell$ to be $\tilde{Q}=\left[\tilde{q}_{1}, \tilde{q}_{3}, \ldots, \tilde{q}_{2 \ell-1}\right]$ where $\tilde{q}_{k}=Q\left(V_{k}\left(\gamma_{\ell}\right)\right)$. This is also called a DCT-II.

The $\tilde{Q}$ vector has values of $Q(X)$ when $X$ is an odd power of $\omega_{\ell}$. Those are the roots of $V_{\ell}$. In the ring of RLPs, we have $Q_{3}(X) \equiv Q_{1}(X) Q_{2}(X)\left(\bmod X^{\ell}+X^{-\ell}\right)$ precisely when the $\tilde{Q}_{3}$ vector is the pointwise product of the $\tilde{Q}_{1}$ and $\tilde{Q}_{2}$ vectors.

This leads to a polynomial multiplication algorithm. Given RLPs $Q_{1}$ and $Q_{2}$ with $\operatorname{deg}\left(Q_{1}\right) \leq 2\left(d_{1}-1\right)$ and $\operatorname{deg}\left(Q_{2}\right) \leq 2\left(d_{2}-1\right)$, choose $\ell \geq d_{1}+d_{2}-1$. Take length $\ell$ DCTs of $Q_{1}$ and $Q_{2}$. Find an RLP $Q_{3}$ of degree at most $2(\ell-1)$ such that $\tilde{Q}_{3}$ is the pointwise product of $\tilde{Q}_{1}$ and $\tilde{Q}_{2}$. The error $Q_{0}(X):=Q_{1}(X) Q_{2}(X)-$ $Q_{3}(X)$ has degree at most $\max \left(2 d_{1}+2 d_{2}-4,2 \ell-2\right)<2 \ell$. But $Q_{0}(X)$ is a multiple of $X^{\ell}+X^{-\ell}$ so $Q_{0}(X)=0$.

We need to transition between $Q$ and $\tilde{Q}$. Steidl and Tasche [?, section 3] do this when the length $\ell$ is a power of 2 , say $\operatorname{deg}(Q)<\ell=2^{r}$, as well as for $\ell=3 \cdot 2^{r}$. If $n_{1}$ and $n_{2}$ are nonnegative integers, let $n_{1} \oplus n_{2}$ denote their bitwise exclusive OR. If $n$ is an (up to) $r$-bit binary integer, let $\operatorname{bitrev}_{r}(n)$ denote its bit reversal. For an odd integer $a$, define $\operatorname{aloc}(a)=\operatorname{bitrev}_{r}(\lfloor a / 4\rfloor \oplus\lfloor a / 2\rfloor)$. Figure ?? sketches how to get $\tilde{Q}$ from $Q$.

```
Input: Array with \(2^{r}\) ring elements.
    \(Q(X)\) starts at index 0 , standard basis.
Outputs:
    Evaluations \(Q\left(V_{a}(\gamma)\right)=Q(X) \bmod \left(X-V_{a}(\gamma)\right)\) at index aloc \((a)\) for odd \(a\).
for \(s\) from 1 to \(r\) do
    \(n:=2^{r-s}\);
// Now array has the remainders \(Q(X) \bmod \left(V_{2 n}(X)-V_{2 a n}(\gamma)\right)\)
// at index \(\operatorname{aloc}(a)\) for odd \(a\) with \(1 \leq a<2^{s}\).
    for \(a\) from 1 to \(2^{s}-1\) by 2 do
        \(b:=a ; c:=2^{s+1}-a ;\)
        \(k:=\operatorname{aloc}(a)\);
        Retrieve \(t_{a}:=Q(X) \bmod \left(V_{2 n}(X)-V_{2 a n}(\gamma)\right)\), which starts at index \(k\).
        \(t_{b}:=t_{a} \bmod \left(V_{n}(X)-V_{b n}(\gamma)\right) ;\)
        \(t_{c}:=t_{a} \bmod \left(V_{n}(X)-V_{c n}(\gamma)\right)\);
        Store \(t_{b}\) and \(t_{c}\), each length \(n\), at indices \(k\) and \(k+n\), respectively.
    end for
end for
```

Fig. 1. Forward DCT-II (without final permutation) of length $2^{r}$

Within the $t_{b}$ and $t_{c}$ computations, we observe
$V_{b n}(\gamma)+V_{c n}(\gamma)=V_{(b n+c n) / 2}(\gamma) V_{(b n-c n) / 2}(\gamma)=V_{2^{r}}(\gamma) V_{(b n-c n) / 2}(\gamma)=0$.
The product of the polynomial moduli is $\left(V_{n}(X)-V_{b n}(\gamma)\right)\left(V_{n}(X)-V_{c n}(\gamma)\right)=$ $V_{n}(X)^{2}-V_{a n}(\gamma)^{2}=V_{2 n}(X)-V_{2 a n}(\gamma)$, which is the modulus used for $t_{a}$.

We can compute $V_{a n}(\gamma)$ for successive $a$ with one ring multiply each since $a$ ranges over an arithmetic progression. Replacing $c_{0}+\sum_{m=1}^{2 n-1} c_{m} V_{m}(X)$ by its two remainders $\left(c_{0} \pm c_{n} V_{a n}(\gamma)\right) \cdot 1+\sum_{m=1}^{n-1}\left(c_{m}-c_{2 n-m} \pm c_{n+m} V_{a n}(\gamma)\right) V_{m}(X)$ takes $n$ multiplications per $a$ and can be done in-place. Each $s$ needs about $2^{s-1}(1+n)=2^{s-1}+2^{r-1}$ multiplies. Many optimizations to the standard FFT apply, such as avoiding multiplies by $\pm 1$.

The pseudocode for an inverse transform can reverse the steps, forming $t_{a}$ from $t_{b}$ and $t_{c}$. The formula $t_{a}:=\left(t_{b}+t_{c}\right) / 2+\left(V_{n}(X) / V_{a n}(\gamma)\right)\left(t_{b}-t_{c}\right) / 2$ divides by $V_{a n}(\gamma)$. Instead assume

$$
\begin{aligned}
& t_{b}=n U_{b n}(\gamma)\left(Q(X) \bmod \left(V_{n}(X)-V_{b n}(\gamma)\right)\right), \\
& t_{c}=n U_{c n}(\gamma)\left(Q(X) \bmod \left(V_{n}(X)-V_{c n}(\gamma)\right)\right)
\end{aligned}
$$

When $s=r$ (so $n=1$ ), the code can scale input remainders by $U_{a}(\gamma)$ so this pattern holds when $n=1$. Given these, the reverse inner loop computes $t_{a}:=V_{n}(X)\left(t_{c}-t_{b}\right)+V_{c n}(\gamma)\left(t_{b}+t_{c}\right)$ to preserve the pattern. At the end, during the backward $s=1$ loop, the only case is $n=2^{r-s}$ and $a=1$. The output $Q(X)$ is scaled by a factor of TBD.

As of November, 2007, the DCT has not been implemented. We plan to have it working before the February, 2008 ANTS deadline.

### 6.3 Multiplying General Polynomials by RLPs

In section ?? we will construct an RLP $h(X)$ which will later be multiplied by various $g(X)$. Normally we use the DCT-II to multiply two RLP's and the DFT to multiply two general polynomials.

A DCT-II of length $\ell / 2$ and a DFT of length $\ell$ use different points of evaluation, so they cannot be combined for multiplication of a polynomial by an RLP.

A DCT-I of length $\ell / 2+1$ can mix its outputs with those of a DFT, each evaluating a polynomial at $\ell$-th roots of unity. We achieve this by computing a full DFT of the RLP (after using $X^{\ell}=1$ to avoid negative exponents). To conserve memory, we store only the $\ell+1$ distinct DFT output coefficients. This exploits $h(\omega)=h\left(\omega^{-1}\right)$ for all $\ell$-th roots of unity $\omega$.

In the scrambled output of a decimation-in-time FTT of length $\ell=2^{r}$, the distinct DFT coefficients $h\left(\omega^{i}\right)$ for $0 \leq i \leq \ell / 2$ are at even indices and index 1. We store only these. For $0<2 i<\ell / 2$, coefficients at index $2 i$ and index $m_{i}-2 i$, where $m_{i}=2^{\left\lfloor\log _{2}(i)\right\rfloor+3}-2^{\left\lfloor\log _{2}(i)\right\rfloor+1}-1$, correspond to $h\left(\omega^{\text {bitrev }_{r}(2 i)}\right)$ and $h\left(\omega^{\ell-\mathrm{bitrev}-r(2 i)}\right)$ and thus are equal. For the pointwise product, we can multiply the FFT coefficients of $g$ at index $2 i$ and $m_{i}-2 i$ by the coefficient of $h$ that was at index $2 i$.

## 7 Computing Coefficients of $f$

Assume the $\mathrm{P}+1$ algorithm. The monic RLP $f(X)$ in (??), with roots $\alpha_{1}^{2 k}$ where $k \in S_{1}$, can be constructed using the decomposition of $S_{1}$. The coefficients of $f$ will always be in the base ring since $P_{1} \in \mathbb{Z} / N \mathbb{Z}$.

For the $\mathrm{P}-1$ algorithm, set $\alpha_{1}=b_{1}$ and $P_{1}=b_{1}+b_{1}^{-1}$. The rest of the construction of $f$ for $\mathrm{P}-1$ is identical to that for $\mathrm{P}+1$.

Assume $S_{1}$ and $S_{2}$ are built as in section ??, say $S_{1}=T_{1}+T_{2}+\cdots+T_{m}$ where each $T_{j}$ has an arithmetic progression of prime length, centered at zero. At least one of these has even cardinality since $s_{1}=\left|S_{1}\right|=\prod_{j}\left|T_{j}\right|$ is even. Renumber the $T_{j}$ so $\left|T_{1}\right|=2$ and $\left|T_{2}\right| \geq\left|T_{3}\right| \geq \cdots \geq\left|T_{m}\right|$.

If $T_{1}=\left\{-k_{1}, k_{1}\right\}$, then initialize $F_{1}(X)=\bar{X}+X^{-1}-\alpha_{1}^{2 k_{1}}-\alpha_{1}^{-2 k_{1}}=$ $X+X^{-1}-V_{2 k_{1}}\left(P_{1}\right)$, a monic RLP in $X$ of degree 2.

Suppose $1 \leq j<m$. Given the coefficients of the monic RLP $F_{j}(X)$ with roots $\alpha_{1}^{2 k_{1}}$ for $k_{1} \in T_{1}+\cdots+T_{j}$, we want to construct

$$
\begin{equation*}
F_{j+1}(X)=\prod_{k_{2} \in T_{j+1}} F_{j}\left(\alpha_{1}^{2 k_{2}} X\right) \tag{9}
\end{equation*}
$$

The set $T_{j+1}$ is assumed to be an arithmetic progression of prime length $t=\left|T_{j+1}\right|$ centered at zero with even common difference $2 k$, say $T_{j+1}=\{(-1-$ $t+2 i) k: 1 \leq i \leq t\}$. On the right of (??), group pairs $\pm k_{2}$ when $k_{2} \neq 0$. We need the coefficients of

$$
F_{j+1}(X)= \begin{cases}F_{j}\left(\alpha_{1}^{-2 k} X\right) F_{j}\left(\alpha_{1}^{2 k} X\right), & \text { if } t=2 \\ F_{j}(X) \prod_{i=1}^{(t-1) / 2}\left(F_{j}\left(\alpha_{1}^{4 k i} X\right) F_{j}\left(\alpha_{1}^{-4 k i} X\right)\right), & \text { if } t \text { is odd }\end{cases}
$$

Let $d=\operatorname{deg}\left(F_{j}\right)$, an even number. The monic input $F_{j}$ has $d / 2$ coefficients in $\mathbb{Z} / N \mathbb{Z}$ (not counting the leading 1). The output $F_{j+1}$ will have $t d / 2=$ $\operatorname{deg}\left(F_{j+1}\right) / 2$ such coefficients.

Products such as $F_{j}\left(\alpha_{1}^{4 k i} X\right) F_{j}\left(\alpha_{1}^{-4 k i} X\right)$ can be formed by the method in section ??, using $d$ coefficients to store each product. The interface can pass $\alpha_{1}^{4 k i}+\alpha_{1}^{-4 k i}=V_{4 k i}\left(P_{1}\right) \in \mathbb{Z} / N \mathbb{Z}$ as a parameter instead of $\alpha_{1}^{ \pm 4 k i}$.

For odd $t$, the algorithm in section ?? forms $(t-1) / 2$ such monic products each with $d$ output coefficients. We still need to multiply by the input $F_{j}$. Overall we store $(d / 2)+\frac{t-1}{2} d=t d / 2$ coefficients. Later these $(t+1) / 2$ monic RLPs can be multiplied in pairs, with products overwriting the inputs, until $F_{j+1}$ (with $t d / 2$ coefficients plus the leading 1 ) is ready.

All polynomial products needed for (??), including those in section ??, have output degree at most $t \operatorname{deg}\left(F_{j}\right)=\operatorname{deg}\left(F_{j+1}\right)$, which divides the final $\operatorname{deg}\left(F_{m}\right)=$ $s_{1}$. If we use multiplications modulo $X^{\ell}-1$, for some $\ell>\operatorname{deg}\left(F_{j+1}\right)$, then the products will be exact. Since $\ell_{\max }>s_{1} \geq \operatorname{deg}\left(F_{j+1}\right)$ and $\ell_{\max }$ is a tolerable convolution length, we can always use $\ell=\ell_{\max }$, but a smaller $\ell$ might be better for a particular product.

### 7.1 Scaling by a Power and its Inverse.

Let $F(X)$ be a monic RLP of even degree $d$, say $F(X)=c_{0}+\sum_{i=1}^{d / 2} c_{i}\left(X^{i}+X^{-i}\right)$, where each $c_{i} \in \mathbb{Z} / N \mathbb{Z}$ and $c_{d / 2}=1$. Given $Q \in \mathbb{Z} / N \mathbb{Z}$, where $Q=\gamma+\gamma^{-1}$ for some unknown $\gamma$, we want the $d$ coefficients (excluding the leading 1) of $F(\gamma X) F\left(\gamma^{-1} X\right) \bmod N$ in place of the $d / 2$ such coefficients of $F$. We are allowed a few scalar temporaries and any storage internal to the polynomial multiplier.

Denote $Y=X+X^{-1}$. Rewrite, while pretending to know $\gamma$,

$$
\begin{aligned}
F(\gamma X) & =c_{0}+\sum_{i=1}^{d / 2} c_{i}\left(\gamma^{i} X^{i}+\gamma^{-i} X^{-i}\right) \\
& =c_{0}+\sum_{i=1}^{d / 2} \frac{c_{i}}{2}\left(\left(\gamma^{i}+\gamma^{-i}\right)\left(X^{i}+X^{-i}\right)+\left(\gamma^{i}-\gamma^{-i}\right)\left(X^{i}-X^{-i}\right)\right) \\
& =c_{0}+\sum_{i=1}^{d / 2} \frac{c_{i}}{2}\left(V_{i}(Q) V_{i}(Y)+\left(\gamma-\gamma^{-1}\right) U_{i}(Q)\left(X-X^{-1}\right) U_{i}(Y)\right)
\end{aligned}
$$

Replace $\gamma$ by $\gamma^{-1}$ and multiply to get

$$
\begin{align*}
F(\gamma X) F\left(\gamma^{-1} X\right) & =G^{2}-\left(\gamma-\gamma^{-1}\right)^{2}\left(X-X^{-1}\right)^{2} H^{2} \\
& =G^{2}-\left(Q^{2}-4\right)\left(X-X^{-1}\right)^{2} H^{2} \tag{10}
\end{align*}
$$

where

$$
G=c_{0}+\sum_{i=1}^{d / 2} c_{i} \frac{V_{i}(Q)}{2} V_{i}(Y), \quad H=\sum_{i=1}^{d / 2} c_{i} \frac{U_{i}(Q)}{2} U_{i}(Y)
$$

This $G$ is a (not necessarily monic) RLP of degree at most $d$ in the standard basis, with coefficients in $\mathbb{Z} / N \mathbb{Z}$. This $H$ is another RLP, of degree at most $d-2$, but using the basis $\left\{U_{i}(Y): 1 \leq i \leq d / 2\right\}$. Starting with the coefficient of $U_{d / 2}(Y)$, we can repeatedly use $U_{j+1}(Y)=V_{j}(Y) U_{1}(Y)+U_{j-1}(Y)=V_{j}(Y)+$ $U_{j-1}(Y)$ for $j>0$, along with $U_{1}(Y)=1$ and $U_{0}(Y)=0$, to convert $H$ to standard basis. This conversion costs $O(d)$ additions in $\mathbb{Z} / N \mathbb{Z}$.

Use (??) and (??) to evaluate $V_{i}(Q) / 2$ and $U_{i}(Q) / 2$ for consecutive $i$ as you evaluate the $d / 2+1$ coefficients of $G$ and the $d / 2$ coefficients of $H$. Using the memory model in section ??, write the standard-basis coefficients of $G$ to one DCT buffer and those of $H$ to the other. Take two forward DCTs, square both, and take the inverse DCTs. Retrieve the $d-1$ coefficients of $H^{2}$ and the $d+1$ coefficients of $G^{2}$ as you finish the (??) computation. Discard the leading 1.

## 8 Multipoint Polynomial Evaluation

We have constructed $f=F_{m}$ in (??). The monic RLP $f(X)$ has degree $s_{1}$, say $f(X)=f_{0}+\sum_{j=1}^{s_{1} / 2} f_{j} \cdot\left(X^{j}+X^{-j}\right)=\sum_{j=-s_{1} / 2}^{s_{1} / 2} f_{j} X^{j}$ where $f_{j}=f_{-j} \in \mathbb{Z} / N \mathbb{Z}$.

Assuming the $\mathrm{P}-1$ method (otherwise see section ??), compute $r=b_{1}^{P} \in$ $\mathbb{Z} / N \mathbb{Z}$. Set $\ell=\ell_{\max }$ and $M=\ell-1-s_{1} / 2$.

Equation (??) needs $\operatorname{gcd}(f(X), N)$ where $X=b_{1}^{2 k_{2}+(2 m+1) P}$, for several consecutive $m$, say $m_{1} \leq m<m_{2}$. By setting $x_{0}=b_{1}^{2 k_{2}+\left(2 m_{1}+1\right) P}$, the arguments to $f$ become $x_{0} b_{1}^{2 m P}=x_{0} r^{2 m}$ for $0 \leq m<m_{2}-m_{1}$. The points of evaluation form a geometric progression with ratio $r^{2}$. We can evaluate these for $0 \leq m<$ $\ell-1-s_{1}$ with one convolution of length $\ell$ and $O(\ell)$ setup cost [?, exercise 8.27].

To be precise, set $h_{j}=r^{-j^{2}} f_{j}$ for $-s_{1} / 2 \leq j \leq s_{1} / 2$. Then $h_{j}=h_{-j}$. Set $h(X)=\sum_{j=-s_{1} / 2}^{s_{1} / 2} h_{j} X^{j}$, an RLP. The construction of $h$ does not reference $x_{0}$ - we reuse $h$ as $x_{0}$ varies.

Let $g_{i}=x_{0}^{M-i} r^{(M-i)^{2}}$ for $0 \leq i \leq \ell-1$ and $g(X)=\sum_{i=0}^{\ell-1} g_{i} X^{i}$.
All nonzero coefficients in $g(X) h(X)$ have exponents from $0-s_{1} / 2$ to ( $\ell-$ $1)+s_{1} / 2$. Suppose $0 \leq m \leq \ell-1-s_{1}$. Then $M-m-\ell=-1-s_{1} / 2-m<-s_{1} / 2$ whereas $M-m+\ell=\left(\ell-1+s_{1} / 2\right)+\left(\ell-s_{1}-m\right)>\ell-1+s_{1} / 2$. The coefficient of $X^{M-m}$ in $g(X) h(X)$, reduced modulo $X^{\ell}-1$, is

$$
\begin{gathered}
\sum_{\substack{0 \leq i \leq \ell-1 \\
-s_{1} / 2 \leq j \leq s_{1} / 2 \\
i+j \equiv M-m(\bmod \ell)}} g_{i} h_{j}=\sum_{\substack{0 \leq i \leq \ell-1 \\
-s_{1} / 2 \leq j \leq s_{1} / 2 \\
i+j=M-m}} g_{i} h_{j}=\sum_{j=-s_{1} / 2}^{s_{1} / 2} g_{M-m-j} h_{j} \\
=\sum_{j=-s_{1} / 2}^{s_{1} / 2} x_{0}^{m+j} r^{(m+j)^{2}} r^{-j^{2}} f_{j}=\sum_{j=-s_{1} / 2}^{s_{1} / 2} x_{0}^{m} r^{m^{2}}\left(x_{0} r^{2 m}\right)^{j} f_{j}=x_{0}^{m} r^{m^{2}} f\left(x_{0} r^{2 m}\right) .
\end{gathered}
$$

Since we want only $\operatorname{gcd}\left(f\left(x_{0} r^{2 m}\right), N\right)$, the $x_{0}^{m} r^{m^{2}}$ factors are harmless.
We can compute successive $g_{\ell-i}$ with two ring multiplications each since the ratios $g_{\ell-1-i} / g_{\ell-i}=x_{0} r^{2 i-s_{1}-1}$ form a geometric progression.

### 8.1 Adaptation for $\mathbf{P}+1$ Algorithm

If we replace $b_{1}$ with $\alpha_{1}$, then $r$ becomes $\alpha_{1}^{P}$, which satisfies $r+r^{-1}=V_{P}\left(P_{1}\right)$. The above algebra evaluates $f$ at powers of $\alpha_{1}$. However $\alpha_{1}, r, h_{j}, x_{0}$, and $g_{i}$ lie in an extension ring.

Arithmetic in the extension ring can use a basis $\{1, \sqrt{\Delta}\}$ where $\Delta=P_{1}^{2}-4$. The element $\alpha_{1}$ maps to $\left(P_{1}+\sqrt{\Delta}\right) / 2$. A product $\left(c_{0}+c_{1} \sqrt{\Delta}\right)\left(d_{0}+d_{1} \sqrt{\Delta}\right)$ where $c_{0}, c_{1}, d_{0}, d_{1} \in \mathbb{Z} / N \mathbb{Z}$ can be done using four base-ring multiplications: $c_{0} d_{0}, c_{1} d_{1},\left(c_{0}+c_{1}\right)\left(d_{0}+d_{1}\right), c_{1} d_{1} \Delta$, plus five base-ring additions.

We define linear transformations $E_{1}, E_{2}$ on $(\mathbb{Z} / N \mathbb{Z})[\sqrt{\Delta}]$ so that $E_{1}\left(c_{0}+\right.$ $\left.c_{1} \sqrt{\Delta}\right)=c_{0}$ and $E_{2}\left(c_{0}+c_{1} \sqrt{\Delta}\right)=c_{1}$ for all $c_{0}, c_{1} \in \mathbb{Z} / N \mathbb{Z}$. Extend $E_{1}$ and $E_{2}$ to polynomials by applying them to each coefficient.

To compute $r^{n^{2}}$ for successive $n$, we use recurrences. We observe

$$
\begin{aligned}
r^{n^{2}} & =r^{(n-1)^{2}+2} \cdot V_{2 n-3}\left(r+r^{-1}\right)-r^{(n-2)^{2}+2} \\
r^{n^{2}+2} & =r^{(n-1)^{2}+2} \cdot V_{2 n-1}\left(r+r^{-1}\right)-r^{(n-2)^{2}}
\end{aligned}
$$

After initializing the variables $\mathrm{r} 1[i]:=r^{i^{2}}, \mathrm{r} 2[i]:=r^{i^{2}+2}, \mathrm{v}[i]:=V_{2 i+1}\left(r+r^{-1}\right)$ for two consecutive $i$, we can compute $\mathrm{r} 1[i]=r^{i^{2}}$ for larger $i$ in sequence by

$$
\begin{align*}
\mathrm{r} 1[i] & :=\mathrm{r} 2[i-1] \cdot \mathrm{v}[i-2]-\mathrm{r} 2[i-2],  \tag{11}\\
\mathrm{r} 2[i] & =\mathrm{r} 2[i-1] \cdot \mathrm{v}[i-1]-\mathrm{r} 1[i-2], \\
\mathrm{v}[i] & :=\mathrm{v}[i-1] \cdot V_{2}(r+1 / r)-\mathrm{v}[i-2] .
\end{align*}
$$

Since we won't use $\mathrm{v}[i-2]$ and $\mathrm{r} 2[i-2]$ again, we can overwrite them with $\mathrm{v}[i]$ and $\mathrm{r} 2[i]$. For the computation of $r^{-n^{2}}$ where $r$ has norm 1 , we can use $r^{-1}$ as input, by taking the conjugate.

All $\mathrm{v}[i]$ are in the base ring but $\mathrm{r} 1[i]$ and $\mathrm{r} 2[i]$ are in the extension ring. Each application of (??) takes five base-ring multiplications (compared to two multiplications per $r^{n^{2}}$ in the $\mathrm{P}-1$ algorithm).

We can compute successive $g_{i}=x_{0}^{M-i} r^{(M-i)^{2}}$ similarly. One solution to (??) is $\mathrm{r} 1[i]=g_{i}, \mathrm{r} 2[i]=r^{2} g_{i}, \mathrm{v}[i]=x_{0} r^{2 M-2 i-1}+x_{0}^{-1} r^{1+2 i-2 M}$. Again each $\mathrm{v}[i]$ is in the base ring, so (??) needs only five base-ring multiplications.

If we try to follow this approach for the multipoint evaluation, we need twice as much space for an element of $(\mathbb{Z} / N \mathbb{Z})[\sqrt{\Delta}]$ as one of $\mathbb{Z} / N \mathbb{Z}$. We also need a convolution routine for the extension ring.

If $p$ divides the coefficient of $X^{M-m}$ in $g(X) h(X)$, then $p$ divides both coordinates thereof. The coefficients of $g(X) h(X)$ occasionally lie in the base ring, making $E_{2}(g(X) h(X))$ a poor choice for the gcd with $N$. Instead we compute

$$
E_{1}(g(X) h(X))=E_{1}(g(X)) E_{1}(h(X))+\Delta E_{2}(g(X)) E_{2}(h(X))
$$

The RLPs $E_{1}(h(X))$ and $E_{2}(\Delta h(X))$ can be computed once and their DCT transforms (convolution length $\ell_{\max }$ ) saved in the two DCT buffers. To compute $E_{2}(\Delta h(X))$, multiply $E_{2}(\mathrm{r} 1[i])$ and $E_{2}(\mathrm{r} 2[i])$ by $\Delta$ after initializing for two consecutive $i$, before applying (??).

Later, as each $g_{i}$ is computed we insert $E_{2}\left(g_{i}\right)$ into a DFT input buffer while saving $E_{1}\left(g_{i}\right)$ for later use. After forming $E_{2}(g(X)) E_{1}(h(X))$, retrieve and save coefficients of $X^{M-m}$ for $0 \leq m \leq \ell-1-s_{1}$. Insert saved $E_{1}\left(g_{i}\right)$ into the (now fresh) DFT input buffer. Form the $E_{1}(g(X)) E_{2}(\Delta h(X))$ product and the sum.

Or we can compute one coordinate of $g_{i}$ at a time, at the cost of computing each $\mathrm{v}[i]$ twice, by applying $E_{1}$ (or $E_{2}$ ) to each line in (??).

## 9 Memory Allocation Model

We aim to fit our major data into the following:
(MZNZ) An array with $s_{1} / 2$ elements of $\mathbb{Z} / N \mathbb{Z}$, for convolution inputs and outputs. This is used during polynomial construction.
This is not needed during $\mathrm{P}-1$ evaluation. During $\mathrm{P}+1$ evaluation, it grows to $\ell_{\max }$ elements of $\mathbb{Z} / N \mathbb{Z}$ (if we compute both coordinate of each $g_{i}$ together, saving one of them), or $\ell_{\max }-s_{1}$ elements (if we compute the coordinates individually).
(MDCT1),
(MDCT2) Two buffers, each holding $\ell_{\max } / 2+1$ outputs from a DCT (convolution length $\left.\ell_{\max }\right)$. Each can instead hold a pointwise product. These are work areas during forward and inverse DCTs.
After the construction of $h$ from $f$ in section ??, the length- $\ell_{\text {max }}$ DFT of $h$ (or of $E_{1}(h)$ and $E_{2}(\Delta h)$ ) is computed once and stored here (using only (MDCT1) for $\mathrm{P}-1$ but both buffers for $\mathrm{P}+1$ ). Then $f$ and $h$ are discarded. The DFT of $h$ evaluates $h(\omega)$ for all $\ell_{\text {max }}$-th roots of unity, but needs only $\left\lfloor\ell_{\max } / 2\right\rfloor+1$ locations due to $h(X)=h(1 / X)$.
(MDFT) A buffer holding $\ell_{\max }$ outputs from a DFT (convolution length $\ell_{\max }$ ). This can also hold a pointwise product. It is a work area during forward and inverse DFTs. This is used during polynomial evaluation, both $P \pm 1$.

During the construction of $F_{j+1}$ from $F_{j}$, if we need to multiply pairs of monic RLPs occupying adjacent locations within (MZNZ) (without the leading 1's), then we can write one DCT input to each DCT buffer. After two forward DCTs, a pointwise product, and an inverse DCT, retrieve all wanted coefficients of the product, overwriting the inputs within (MZNZ).

During polynomial evaluation for $\mathrm{P}-1$, we need only (MDCT1) and (MDFT). Send each $g_{i}$ coefficient to (MDFT) as $g_{i}$ is computed. When (MDFT) fills (with $\ell_{\text {max }}$ entries), do a length- $\ell_{\max }$ forward DFT on (MDFT), pointwise multiply by the saved DCT output from $h$ in (MDCT1), and do an inverse DFT in (MDFT). Retrieve each polynomial, compute their product, and take a GCD with $N$.

### 9.1 Potentially Large $\boldsymbol{B}_{\mathbf{2}}$

Nowadays (2007) a typical PC memory is 4 gigabytes. The median size of composite cofactors $N$ in the Cunningham project http://homes.cerias.purdue.

Table 1. Estimated memory usage (quadwords) while factoring 230-digit number.

| Array <br> name | Construct $f$. <br> Both $P \pm 1$ | Evaluate $f$. <br> $\mathrm{P}-1$ | Evaluate $f$. <br> $\mathrm{P}+1$ |
| :--- | ---: | ---: | ---: |
| (MZNZ) | $12 s_{1} / 2$ |  | 0 |
| $\left(\begin{array}{l}12 \ell_{\max } \\ \left(\text { or } 12\left(\ell_{\max }-s_{1}\right)\right)\end{array}\right.$ |  |  |  |
| (MDCT1) | $25\left(1+\ell_{\max } / 2\right)$ | $25\left(1+\ell_{\max } / 2\right)$ | $25\left(1+\ell_{\max } / 2\right)$ |
| (MDCT2) | $25\left(1+\ell_{\max } / 2\right)$ | 0 | $25\left(1+\ell_{\max } / 2\right)$ |
| (MDFT) | 0 | $25 \ell_{\max }$ | $25 \ell_{\max }$ |
| Totals, if <br> $s_{1}=\ell_{\max } / 2$ | $28 \ell_{\max }+O(1)$ | $37.5 \ell_{\max }+O(1)$ | $62 \ell_{\max }+O(1)$ <br> (or $\left.56 \ell_{\max }+O(1)\right)$ |

edu/~ssw/cun/index.html is about 230 decimal digits, which fits in twelve 64 -bit words (called quadwords). Table ?? estimates the memory requirements during stage 2 , when factoring a 230 -digit number, for both polynomial construction and polynomial evaluation phases, assuming convolutions use the NTT approach in section ??. The product of our NTT prime moduli must be at least $\ell_{\max }(N-1)^{2}$. If $N$ is below $0.99 \cdot\left(2^{63}\right)^{25} \approx 10^{474}$, then it will suffice to have 25 NTT primes, each 63 bits.

The $\mathrm{P}-1$ polynomial construction phase uses an estimated $28 \ell_{\max }$ quadwords, vs. $37.5 \ell_{\text {max }}$ quadwords during polynomial evaluation. Four gigabytes is 537 million quadwords. A possible value is $\ell_{\max }=2^{23}$, which needs 315 million quadwords. When transform length $3 \cdot 2^{k}$ is supported, we could use $\ell \max =3 \cdot 2^{22}$ which needs 472 million quadwords.

We might use $P=3 \cdot 5 \cdot 7 \cdot 11 \cdot 13 \cdot 17 \cdot 19 \cdot 23=111546435$, for which $\phi(P)=36495360=2^{13} \cdot 3^{4} \cdot 5 \cdot 11$. We choose $s_{2} \mid \phi(P)$ so that $s_{2}$ is close to $\phi(P) /\left(\ell_{\max } / 2\right) \approx 8.7$, i.e. $s_{2}=9$ and $s_{1}=4055040$, giving $s_{1} / \ell_{\max } \approx 0.48$.

We can do 9 convolutions, one for each $k_{2} \in S_{2}$. We will be able to find $p \mid N$ if $b_{1}^{q} \equiv 1(\bmod p)$ where $q$ satisfies (??) with $m<\ell_{\max }-s_{1}=4333568$. As described in section ??, the effective value of $B_{2}$ will be about $9.66 \cdot 10^{14}$.

## 10 Opportunities for Parallelization

Modern PC's are multi-core, typically with 2-4 CPUs (cores) and a shared memory. When running on such systems, it is desirable to utilize multiple cores.

While building $h(X)$ and $g(X)$ in section ??, each core can process a contiguous block of subscripts. Use the explicit formulas to compute $r^{-j^{2}}$ or $g_{i}$ for the first two elements of a block, and the recurrences elsewhere.

If convolutions use NTT's and the number of processors divides the number of primes, then allocate the primes evenly across the processors. The DCT and DFT buffers in section ?? can have separate subbuffers for each prime. On NUMA architectures, the memory for each subbuffer should be allocated locally to the processor that will process it. Accesses to remote memory occur
only when converting the $h_{j}$ and $g_{i}$ to residues modulo small primes, and when reconstructing the coefficients of $g(x) h(x)$ with the CRT.

## 11 Our Implementation

Our implementation is based on GMP-ECM, an implementation of $\mathrm{P}-1, \mathrm{P}+1$ and the Elliptic Curve Method for integer factorization. It uses the GMP library for arbitrary precision arithmetic. The code for stage 1 of $\mathrm{P}-1$ and $\mathrm{P}+1$ is unchanged; the code for the new stage 2 has been written from scratch and will replace the previous implementation [?] which used product trees of cost $O\left(n(\log n)^{2}\right)$ modular multiplications for building polynomials of degree $n$ and a variant of Montgomery's POLYEVAL [?] algorithm for multipoint evaluation which has cost $O\left(n(\log n)^{2}\right)$ modular multiplications and $O(n \log n)$ memory. The practical limit for $B_{2}$ was about $10^{14}-10^{15}$.

GMP-ECM includes modular arithmetic routines, using e.g. Montgomery's REDC [?] or fast reduction modulo number of the form $2^{n} \pm 1$. It also includes routines for polynomial arithmetic, in particular convolution products. One algorithm available for this purpose is a small prime NTT/CRT (but without the DCT variant). Its current implementation allows only for power-of-two transform lengths. Another is Kronecker-Schönhage's segmentation method [?], which is faster than the NTT if the modulus is large and the convolution length is comparatively small, and it works for any convolution length. Its main disadvantage is significantly higher memory use, reducing the possible convolution length.

On a 2.4 GHz Opteron with 8 GB memory, $\mathrm{P}-1$ stage 2 on a 250 -digit composite number with $B_{2}=1.2 \cdot 10^{15}$, using the NTT for the convolution, can use $P=64579515, \ell_{\max }=2^{24}, s_{1}=7434240, s_{2}=3$ and takes 46 minutes.

On the same machine, $\mathrm{P}+1$ stage 2 on a 243 -digit number with $B_{2}=10^{15}$ can use $P=111546435, l=2^{23}, s_{1}=3649536, s_{2}=10$ and takes 86 minutes.

## 12 Some Results

We ran at least one of $P \pm 1$ on over 1500 composite cofactors, including
(a) Richard Brent's tables with $b^{n} \pm 1$ factorizations for $13 \leq b \leq 99$;
(b) Fibonacci and Lucas numbers $F_{n}$ and $L_{n}$ with $n<2000$, or $n<10000$ and cofactor size $<10^{300}$;
(c) Cunningham cofactors of $12^{n} \pm 1$ with $n<300$.

The $B_{1}$ and $B_{2}$ values varied, with $10^{11}$ and $10^{15}$ being typical. Table ?? has new large prime factors $p$ and the largest factors of the corresponding $p \pm 1$.

The 52 -digit factor of $47^{146}+1$ and the 60 -digit factor of $L_{2366}$ each set a new record for the $\mathrm{P}+1$ factoring algorithm upon their discovery. The previous record was a 48-digit factor of $L_{1849}$, found by the second author in March 2003.

The 49-digit factor of $75^{128}+1$ has $q=3288338823576187$, a 16 -digit prime. To our knowledge, this is the largest prime in the group order associated with any factor found by the $\mathrm{P}-1, \mathrm{P}+1$ or Elliptic Curve methods of factorization.

Table 2. Large $P \pm 1$ factors found

| Input <br> Method | Factor $p$ found <br> Largest factors of $p \pm 1$ | Size |
| :---: | :---: | :---: |
| $68^{118}+1$ | $7506686348037740621097710183200476580505073749325089^{*}$ | c 151 |
| $\mathrm{P}-1$ | $22807 \cdot 480587 \cdot 14334767 \cdot 89294369 \cdot 4649376803 \cdot 5380282339$ | p 52 |
| $73^{\mathrm{P}-9}-1$ | 76227040047863715568322367158695720006439518152299 | c 191 |
| $\mathrm{P}-1$ | $12491 \cdot 37987 \cdot 156059 \cdot 2244509 \cdot 462832247372839$ | p 50 |
| $75^{128}+1$ | 2180637877078565656090569549812536273194443435521 | c 183 |
| $\mathrm{P}-1$ | $37489 \cdot 19791649 \cdot 162960727 \cdot 3288338823576187$ | p 49 |
| $47^{146}+1$ | 7986478866035822988220162978874631335274957495008401 | c 235 |
| $\mathrm{P}+1$ | $20540953 \cdot 56417663 \cdot 1231471331 \cdot 1632221953 \cdot 843497917739$ | p 52 |
| $L_{2366}$ | 725516237739635905037132916171116034279215026146021770250523 | c 290 |
| $\mathrm{P}+1$ | $932677 \cdot 62754121 \cdot 1988258341 \cdot 751245344783 \cdot 483576618980159$ | p 60 |

The largest $q$ reported in Table 2 of [?] is $q=6496749983$ ( 10 digits), for a 19-digit factor $p$ of $2^{895}+1$. That table includes a 34 -digit factor of the Fibonacci number $F_{575}$, which was the $\mathrm{P}-1$ record in 1989.

The largest $\mathrm{P}-1$ factor reported in [?, pp. 538-539] is a 58 -digit factor of $2^{2098}+1$ with $q=9909876848747$ (13 digits). Site http://www.loria.fr/ ~zimmerma/records/Pminus1.html has other records, including a 66-digit factor of $960^{119}-1$ found by $\mathrm{P}-1$ for which $q=2110402817$ (only ten digits).

## 13 Acknowledgements

We thank Paul Zimmermann for his advice and guidance.
We thank Centrum voor Wiskunde en Informatica (CWI, Amsterdam) and INRIA for providing huge amounts of computer time for this work.

## References

1. Alfred V. Aho, John E. Hopcroft, and Jeffrey D. Ullman, The Design and Analysis of Computer Algorithms, Addison-Wesley, 1974.
2. Peter L. Montgomery, Modular Multiplication without Trial Division, Math. Comp., v. 44(1985), pp. 519-521
3. Peter L. Montgomery, Speeding the Pollard and Elliptic Curve Methods of Factorization, Math. Comp., v. 48(1987), pp. 243-264.
4. Peter L. Montgomery and Robert D. Silverman, An FFT Extension to the P-1 Factoring Algorithm, Math. Comp., v. 54(1990), pp. 839-854.
5. Peter L. Montgomery. An FFT Extension to the Elliptic Curve Method of Factorization. UCLA dissertation, 1992. Available at ftp://ftp.cwi.nl/pub/pmontgom.
6. J.M. Pollard, Theorems on factorization and primality testing, Proc. Cambridge Philosophical Society, volume 76(1974), pp. 521-528.
7. G. Steidl and M. Tasche, "A polynomial approach to fast algorithms for discrete Fourier-cosine and Fourier-sine transforms", Math. Comp., v. 56(1991), pp. 281296.
8. H. C. Williams, A $p+1$ Method of Factoring, Math. Comp., v. 39(1982), pp. 225234.
9. Paul Zimmermann and Bruce Dodson. 20 Years of ECM, In F. Hess, S. Pauli, M. Pohst (Eds.), Algorithmic Number Theory, 7th International Symposium, ANTSVII volume 4076 of Lecture Notes in Computer Science, pp. 525-542. SpringerVerlag, 2006.
