ON TESTING THE DIVISIBILITY OF LACUNARY POLYNOMIALS BY CYCLOTOMIC POLYNOMIALS

Michael Filaseta* and Andrzej Schinzel

1. Introduction and the Main Theorems

This note describes an algorithm for determining whether a given polynomial $f(x) \in \mathbb{Z}[x]$ has a cyclotomic divisor. In particular, the algorithm works well when the number of non-zero terms is small compared to the degree of $f(x)$. This work is based on papers of H. B. Mann [3] and J. H. Conway and A. J. Jones [1].

For m a positive integer, we define $\Phi_m(x)$ to be the mth cyclotomic polynomial, and let $\zeta_m = e^{2\pi i/m}$. If $m = k\ell$ where k and ℓ are relatively prime positive integers, then it is not difficult to see that $\mathbb{Q}(\zeta_m) = \mathbb{Q}(\zeta_k, \zeta_\ell)$. Letting ϕ denote Euler's ϕ-function, we observe that $[\mathbb{Q}(\zeta_k) : \mathbb{Q}] = \phi(k)$ and $[\mathbb{Q}(\zeta_m) : \mathbb{Q}] = \phi(m) = \phi(k)\phi(\ell)$. It follows that $[\mathbb{Q}(\zeta_m) : \mathbb{Q}(\zeta_k)] = \phi(\ell)$ and that $\Phi_\ell(x)$ is the minimal polynomial for ζ_ℓ over the field $\mathbb{Q}(\zeta_k)$. In particular, $\{1, \zeta_\ell, \zeta_\ell^2, \dots, \zeta_\ell^{\phi(\ell)-1}\}$ forms a basis for $\mathbb{Q}(\zeta_m)$ over $\mathbb{Q}(\zeta_k)$. We will mainly be interested in the case when ℓ is a prime power. For integers a, b, and m with $m > 0$, we write $a \equiv b \pmod m$ if $m|(a-b)$ and use the notation $a \bmod m$ to represent the unique $b \equiv a \pmod m$ such that $0 \le b < m$. For $f(x)$, $g(x)$, and $w(x)$ in $\mathbb{Q}[x]$ with $\deg w(x) \ge 1$, we write $f(x) \equiv g(x) \pmod{w(x)}$ if $w(x)|(f(x) - g(x))$, and we use the notation $f(x) \bmod w(x)$ to denote the unique polynomial $g(x) \equiv f(x) \pmod{w(x)}$ with either $g(x) \equiv 0$ or $0 \le \deg g(x) < \deg w(x)$.

Theorem 1. *Let $f(x) \in \mathbb{Z}[x]$. Let m, k, and ℓ be positive integers, with k and ℓ relatively prime, $m = k\ell$, and $\Phi_m(x)$ dividing $f(x)$. Let*

$$f(xy) \bmod \Phi_\ell(y) = \sum_{j=0}^{\phi(\ell)-1} a_j(x) y^j. \tag{1}$$

Then $\Phi_k(x)$ divides each $a_j(x)$ for $0 \le j \le \phi(\ell) - 1$.

Proof. Observe that $m = k\ell$ and k and ℓ being relatively prime imply that $\zeta_k\zeta_\ell$ is a primitive mth root of unity. Taking $x = \zeta_k$ and $y = \zeta_\ell$, we deduce that

$$\sum_{j=0}^{\phi(\ell)-1} a_j(\zeta_k)\zeta_\ell^j = 0.$$

On the other hand, $\{1, \zeta_\ell, \zeta_\ell^2, \dots, \zeta_\ell^{\phi(\ell)-1}\}$ is a basis for $\mathbb{Q}(\zeta_m)$ over $\mathbb{Q}(\zeta_k)$. Hence, $a_j(\zeta_k) = 0$ for $0 \le j \le \phi(\ell) - 1$, and the result follows. ■

*The first author was supported by NSF Grant DMS-9400937 and NSA Grant MDA904-97-1-0035.

Corollary 1. *Let $f(x) \in \mathbb{Z}[x]$. Let m be a positive integer for which $\Phi_m(x)$ divides $f(x)$. Let*

$$f(xy) \bmod \Phi_m(y) = \sum_{j=0}^{\phi(m)-1} a_j(x) y^j.$$

Then $a_j(1) = 0$ for $0 \le j \le \phi(m) - 1$.

The Corollary immediately follows by considering the case $\ell = m$ and $k = 1$ in Theorem 1. We note that the converses of both Theorem 1 and this Corollary hold. In other words, if the $a_j(x)$ are defined as in Theorem 1 and if $\Phi_k(x)$ divides each $a_j(x)$, then $\Phi_m(x)$ divides $f(x)$ (provided k and ℓ are relatively prime); this follows along lines similar to the argument given for Theorem 1. Also, we note that the polynomials $a_j(x)$ defined in Theorem 1 will necessarily have integer coefficients since $\Phi_\ell(y)$ is monic. In addition, we observe that if $\ell = 1$ in Theorem 1, then the sum on the right side of (1) consists of one term, namely $a_0(x) = f(x)$.

Theorem 2. *Let $f(x) \in \mathbb{Z}[x]$ have N non-zero terms. Suppose n is a positive integer such that $\Phi_n(x) | f(x)$. Suppose further that $p_1, p_2, \ldots, p_k$ are distinct primes satisfying*

$$2 + \sum_{j=1}^{k} (p_j - 2) > N.$$

Let e_j be the non-negative integer for which $p_j^{e_j} || n$. Then for at least one $j \in \{1, 2, \ldots, k\}$, we have $\Phi_m(x) | f(x)$ where $m = n/p_j^{e_j}$.

Proof. We first describe and then make use of Theorem 5 from [1]. For r a positive integer, define $\gamma(r) = 2 + \sum_{p|r} (p-2)$. Following [1], we call a vanishing sum S minimal if no proper subsum of S vanishes. We will be interested in sums $S = \sum_{j=1}^{t} a_j \omega_j$ where t is a positive integer, each a_j is a non-zero rational number and each ω_j is a root of unity. We refer to the reduced exponent of such an S as the least positive integer r for which $(\omega_i/\omega_1)^r = 1$ for all $i \in \{1, 2, \ldots, t\}$. Theorem 5 of [1] asserts then that if $S = \sum_{j=1}^{t} a_j \omega_j$ is a minimal vanishing sum, then $t \ge \gamma(r)$ where r is the reduced exponent of S. (Also, note that Theorem 1 of [1] implies that the reduced exponent r of a minimal vanishing sum is necessarily squarefree.)

To prove Theorem 2, we suppose as we may that $e_j > 0$ for each $j \in \{1, 2, \ldots, k\}$. We write $f(x) = \sum_{i=1}^{s} f_i(x)$ where each $f_i(x)$ is a non-zero polynomial divisible by $\Phi_n(x)$, no two $f_i(x)$ have terms involving x to the same power, and s is maximal. Thus, each $f_i(\zeta_n)$ is a minimal vanishing sum. For each $i \in \{1, 2, \ldots, s\}$, we write $f_i(x) = x^{b_i} g_i(x^{d_i})$ where b_i and d_i are non-negative integers chosen so that $g_i(0) \neq 0$ and the greatest common divisor of the exponents appearing in $g_i(x)$ is 1. Then $g_i(\zeta_n^{d_i})$ is a minimal vanishing sum with reduced exponent $n/\gcd(n, d_i)$. If t_i denotes the number of non-zero terms of $g_i(x)$, we deduce from Theorem 5 of [1] that

$$N = \sum_{i=1}^{s} t_i \ge \sum_{i=1}^{s} \gamma\left(\frac{n}{\gcd(n, d_i)}\right) \ge 2s + \sum_{j=1}^{k} (p_j - 2) \left| \left\{ 1 \le i \le s : p_j \text{ divides } \frac{n}{\gcd(n, d_i)} \right\} \right|.$$

The inequality in Theorem 2 implies that at least one of the expressions $|\{1 \le i \le s : p_j|(n/\gcd(n,d_i))\}|$ is zero. In other words, for some $j \in \{1,2,\ldots,k\}$ and every $i \in \{1,2,\ldots,s\}$, we have $p_j^{e_j}|d_i$. Setting $m = n/p_j^{e_j}$ and $d_i' = d_i/p_j^{e_j}$, we obtain that $g_i(\zeta_m^{d_i'}) = 0$. Since $\gcd(m,p_j) = 1$, $\zeta_m^{p_j^{e_j}}$ is a primitive mth root of unity and we deduce $g_i(\zeta_m^{d_i}) = 0$. As this is true for every $i \in \{1,2,\ldots,s\}$, we conclude $f(\zeta_m) = 0$, establishing the theorem. ■

Corollary 2. *Let $f(x) \in \mathbb{Z}[x]$ have N non-zero terms. If $f(x)$ is divisible by a cyclotomic polynomial, then there is a positive integer m such that $2 + \sum_{p|m}(p-2) \le N$ and $\Phi_m(x)|f(x)$.*

The above is a direct consequence of Theorem 2. Observe that it follows easily from Corollary 2 that if $f(x)$ is divisible by a cyclotomic polynomial, then there is a positive integer m such that every prime divisor of m is $\le N$ and $\Phi_m(x)|f(x)$.

2. The Algorithm

We now describe how to determine if a given $f(x) \in \mathbb{Z}[x]$ has a cyclotomic divisor. The algorithm we describe has been implemented using MAPLE V, Release 5, and an interactive World Wide Web page that allows users to input polynomials $f(x)$ and make use of the algorithm is available at the web address

`http://www.math.sc.edu/~filaseta/cyclotomic.html`.

The web page was developed jointly by Douglas Meade and the first author. A copy of the MAPLE program can also be downloaded there (but access to MAPLE is necessary to take advantage of the downloaded program). In particular, we note that on a Sun Ultra 1, the program handles a random polynomial with 100 non-zero terms, with coefficients bounded by 1000000 and with degree 1000000 in approximately two and a half minutes. It is also possible to run the program with polynomials of degree up to 10^{100}.

Corollary 2 implies that we need only consider the possibility that $\Phi_m(x)|f(x)$ where each prime divisor of m is $\le N$, the number of non-zero terms of $f(x)$. For each prime $p \le N$, we compute the value of

$$r(p) = \left[\frac{\log \deg f}{\log p}\right] + 1.$$

Observe that if $p^e|m$, then

$$\deg f(x) \ge \deg \Phi_m(x) = \phi(m) \ge \phi(p^e) = p^{e-1}(p-1) \ge p^{e-1}$$

so that $e \le r(p)$.

When computing the $a_j(x)$ associated with Theorem 1 or determining whether $f(x)$ is divisible by a cyclotomic polynomial $\Phi_\ell(x)$ in this section, we reduce the exponents of $f(x)$ modulo ℓ (as in [2]) and then perform a division by $\Phi_\ell(x)$. To be more precise, suppose

$$f(x) = \sum_{j=0}^{r} b_j x^{d_j} \tag{2}$$

where the b_j denote non-zero integers and the d_j denote distinct non-negative integers. To compute $f(xy) \bmod \Phi_\ell(y)$, one can first compute the values of $\overline{d}_j = d_j \bmod \ell$ to obtain

$$f(xy) \equiv \sum_{j=0}^{r} b_j x^{d_j} y^{\overline{d}_j} \equiv \sum_{i=0}^{\ell-1} c_i(x) y^i \pmod{\Phi_\ell(y)}, \tag{3}$$

where the $c_i(x)$ are obtained by combining equal values of $\overline{d}_j$. Note that this definition of the $c_i(x)$ implies that $f(x) = \sum_{j=0}^{\ell-1} c_i(x)$. Now, each of $y^{\overline{d}_j} \bmod \Phi_\ell(y)$ can be computed directly. We will be interested in the case that $\ell = p^e$ for some prime p and some positive integer e. If $\overline{d}_j < \phi(\ell) = p^{e-1}(p-1)$, then the value of $y^{\overline{d}_j}$ will remain unchanged when we consider it mod $\Phi_\ell(y)$. If $\overline{d}_j \geq p^{e-1}(p-1)$, then we use the reduction

$$y^{\overline{d}_j} \bmod \Phi_\ell(y) = -\sum_{u=1}^{p-1} y^{\overline{d}_j - p^{e-1}u}, \tag{4}$$

which follows from $\Phi_\ell(y) = \Phi_p(y^{p^{e-1}}) = \sum_{j=0}^{p-1} y^{p^{e-1}j}$. Observe that since $\overline{d}_j < p^e$, the exponents appearing on the right of (4) are in the interval $[0, p^{e-1}(p-1))$. Also, the exponents on the right of (4) are different for different choices of $\overline{d}_j$ in $[p^{e-1}(p-1), p^e)$ (as the integers in this interval are incongruent modulo p^{e-1}).

A simple algorithm for testing $f(x)$ for divisibility by a cyclotomic polynomial can be described as follows. Let $p_1, p_2, \ldots, p_r$ denote the complete list of primes $\leq N$ with $p_1 < p_2 < \cdots < p_r$. Consider each integer

$$m = \prod_{j=1}^{r} p_j^{e_j} \qquad \text{where } 0 \leq e_j \leq r(p_j), \tag{5}$$

and check directly if $\Phi_m(x) | f(x)$. If some $\Phi_m(x)$ divides $f(x)$, then the algorithm stops and indicates so. Otherwise, the algorithm outputs that $f(x)$ has no cyclotomic divisor. This simple algorithm is not, however, as efficient as we would like. In particular, to handle polynomials as described in the first paragraph of this section would require testing over 6×10^{19} different values of m. Some improvement can be made by considering only m for which $\phi(m) \leq \deg f(x)$, but our main improvement will be of a different nature.

Given a positive integer ℓ, consider $a_j(x)$ as defined in (1). Observe that if some $a_j(x)$ has only one non-zero term, then it cannot be divisible by $\Phi_k(x)$ for any positive integer k. It follows then from Theorem 1 that $\Phi_{k\ell}(x)$ does not divide $f(x)$ for any positive integer k. We use this information to reduce the number of m we need consider in (5). We now describe how this is achieved.

For $1 \leq j \leq r$, define $B_j = r(p_j)$. Let

$$S = \{(e_1, e_2, \ldots, e_r) : 0 \leq e_j \leq B_j \text{ for } 1 \leq j \leq r\}.$$

Thus, S consists of the r-tuples formed from the exponents appearing in (5). Define a lexicographical ordering on the elements of S. More precisely, if $A = (a_1, \ldots, a_r)$ and

$B = (b_1, \ldots, b_r)$ are elements of S, then we say $A < B$ if there is some $i \in \{0, 1, \ldots, r-1\}$ such that $a_1 = b_1, a_2 = b_2, \ldots, a_i = b_i$, and $a_{i+1} < b_{i+1}$. If d and m are positive integers, we say that d exactly divides m and write $d||m$ if $d|m$ and $\gcd(d, m/d) = 1$. Finally, we say that d is obtained from an element $(e_1, e_2, \ldots, e_r)$ of S if $d = p_1^{e_1} \cdots p_r^{e_r}$.

We describe a subroutine `shrinkf` that makes use of Theorem 1. The input for `shrinkf` consists of a polynomial $g(x)$, a prime p, and a positive integer r. We consider $g(x)$ in place of $f(x)$ and $\ell = p^r$ in Theorem 1. The subroutine `shrinkf` computes the $a_j(x)$ given there and outputs a non-zero value of $a_j(x)$ which has as few non-zero terms as possible. The idea is that if we wish to know if there is an m exactly divisible by p^r for which $\Phi_m(x)$ divides $g(x)$, then `shrinkf`(g, p, r) will output a polynomial $a(x)$ which "typically" has less (and never has more) non-zero terms than $g(x)$ and which is divisible by $\Phi_{m/p^r}(x)$.

Suppose d is obtained from $(e_1, e_2, \ldots, e_r)$ in S. Write $d = q_1^{e_1'} \cdots q_s^{e_s'}$ where each q_j is prime with $q_1 < \cdots < q_s \leq N$ and where each e_j' is a positive integer. Setting $f_0(x) = f(x)$, we define recursively the polynomials

$$f_j(x) = \texttt{shrinkf}(f_{j-1}, q_j, e_j') \qquad \text{for } 1 \leq j \leq s.$$

If $\Phi_m(x)|f(x)$ and $d||m$, then Theorem 1 implies inductively that $\Phi_{m/\left(q_1^{e_1'} \cdots q_j^{e_j'}\right)}(x)$ divides $f_j(x)$. In particular, in this case, we have $\Phi_{m/d}(x)|f_s(x)$. Thus, if $f_s(x)$ is a polynomial with exactly one non-zero term and $\Phi_m(x)|f(x)$, then it follows that d cannot exactly divide m. We also observe that if $\Phi_d(x)|f(x)$, then $f_s(1) = 0$.

We wish to determine if $f(x)$ is divisible by $\Phi_m(x)$ for some integer m obtained from an element of S. We go through the elements of S in increasing order as follows. We begin with the smallest element $(0, 0, \ldots, 0)$. We let d denote the integer obtained from the element of S under consideration; initially then $d = 1$. For each element $(e_1, \ldots, e_r)$ of S under consideration and the corresponding d obtained from it, we proceed by computing $f_s(x)$ as defined above (in the case of $d = 1$, we interpret $f_s(x)$ to be $f_0(x) = f(x)$). If $f_s(x)$ is a polynomial with exactly one non-zero term, then we know that $f(x)$ cannot be divisible by $\Phi_m(x)$ for any integer m exactly divisible by d. This is where a savings over the simple algorithm described earlier in this section occurs; we can now ignore integers m obtained from elements of S if $d||m$. To explain how this is done, we let i and j be integers with $0 \leq i \leq j \leq r$ satisfying

$$0 \leq e_i < B_i,\ e_{i+1} = B_{i+1}, e_{i+2} = B_{i+2}, \ldots, e_j = B_j,\ e_{j+1} = 0, e_{j+2} = 0, \ldots, e_r = 0.$$

Here, if a subscript is not in the range $[1, r]$, then it is to be ignored; and if $j = i$, then we take this to mean that the middle equations $e_t = B_t$ with $i+1 \leq t \leq j$ are not present. The case $d = 1$ we equate with $i = j = 0$. If $i \geq 1$, then the next element of S we consider is $(e_1, \ldots, e_{i-1}, e_i + 1, 0, 0, \ldots, 0)$. Given the ordering of the elements of S described earlier, if m is obtained from an element of S between $(e_1, \ldots, e_r)$ and $(e_1, \ldots, e_{i-1}, e_i + 1, 0, \ldots, 0)$, then $d||m$. This implies that we have skipped over considering elements of S which we know correspond to cyclotomic polynomials that cannot divide $f(x)$. If $i = 0$ above, then we stop the algorithm and indicate that $f(x)$ has no cyclotomic divisors. In this case, the remaining elements of S correspond to integers m for which $d||m$ and so they need not be considered. This describes how we proceed if $f_s(x)$ has exactly one non-zero term.

Suppose now that $(e_1, \ldots, e_r)$ and d are as before but $f_s(x)$ has more than one non-zero term. We check if $f_s(1) = 0$. If it does, then it might happen that $\Phi_d(x)|f(x)$ and we determine whether this is the case by a direct computation. If we determine $\Phi_d(x)|f(x)$, then we stop and output that $f(x)$ is divisible by a cyclotomic polynomial. Otherwise, either we have $f_s(1) = 0$ and $\Phi_d(x) \nmid f(x)$ or we have $f_s(1) \neq 0$ (and $\Phi_d(x) \nmid f(x)$ here as well). Note that we could have simply checked if $\Phi_d(x)|f(x)$ without checking the value of $f_s(1)$; the idea behind checking if $f_s(1) = 0$, however, is that we expect usually $f_s(1) \neq 0$ and checking the value of $f_s(1)$ saves time (i.e., division by $\Phi_d(x)$ is more time consuming). If $\Phi_d(x) \nmid f(x)$, then we simply consider the next element of S greater than $(e_1, \ldots, e_r)$ in our ordering of the elements of S. (Some savings is gained by observing that the $f_s(x)$ associated with each subsequent d can be computed by a single call to `shrinkf` by making use of previously computed $f_j(x)$.)

The above describes the algorithm. To justify the algorithm works is relatively simple. First, observe that if the algorithm indicates $f(x)$ has a cyclotomic divisor, it does as this indication will be as a result of a direct check of a possible cyclotomic divisor as in the paragraph above. On the other hand, for each m as in (5), either the algorithm checks if $\Phi_m(x)|f(x)$ when $d = m$ (directly or by computing $f_s(1)$) or the algorithm has skipped over considering $\Phi_m(x)$ as a divisor of $f(x)$ because the algorithm has verified that $\Phi_m(x) \nmid f(x)$ through the use of some exact divisor d of m. Thus, the algorithm determines if there is an m as in (5) with $\Phi_m(x)$ dividing $f(x)$, and this is sufficient for deciding whether $f(x)$ is divisible by a cyclotomic polynomial.

3. Running Times and Other Concluding Remarks

One difficulty in estimating the running time of the algorithm in Section 2 is in finding good bounds for the total number of elements of S that are skipped over as we determine d for which $\Phi_m(x)$ does not divide $f(x)$ for positive integers m exactly divisible by d. In addition, we checked whether $\Phi_m(x)$ divides $f(x)$ by a direct division. Even after reducing the exponents of $f(x)$ modulo m, the division can cost considerable time as the size of m can be large and can even exceed $\deg f(x)$. To get a good theoretical bound for the maximal running time of an algorithm which determines whether $f(x)$ has a cyclotomic factor, we make the following changes in the algorithm just described.

We alter `shrinkf` so that, given input $g(x)$, p, and r as before, `shrinkf` returns not just one but every $a_j(x)$ obtained from Theorem 1 (discarding multiplicities). If there are t distinct non-zero $a_j(x)$, then the output is a t-tuple with components consisting of these $a_j(x)$ (in any order).

We go through the elements of S using their lexicographical ordering as before, checking to see if $\Phi_m(x)$ divides $f(x)$ for some m obtained from an element of S. Let d be a number we are considering, obtained from $(e_1, e_2, \ldots, e_r)$ in S. In particular, we may suppose that at this point in the algorithm, we have already determined that each d' obtained from a previous element of S is such that $\Phi_{d'}(x)$ does not divide $f(x)$ (otherwise the algorithm would have already terminated indicating that $f(x)$ has a cyclotomic divisor). We write $d = q_1^{e_1'} \cdots q_s^{e_s'}$ where, as before, each q_j is prime with $q_1 < \cdots < q_s \leq N$ and each e_j' is a

positive integer. We check initially whether

$$2 + \sum_{j=1}^{s} (q_j - 2) \le N. \tag{6}$$

If not, then Corollary 2 implies that we need not consider any m obtained from an element of S that is divisible by d. As in the previous section, we can skip over those elements of S exactly divisible by d (and, in fact, more).

We consider next the case where (6) holds. Define a positive integer t_1 and polynomials $f_1^{(1)}(x), \ldots, f_1^{(t_1)}(x)$ by

$$\big(f_1^{(1)}(x), \ldots, f_1^{(t_1)}(x)\big) = \texttt{shrinkf}(f, q_1, e_1').$$

Next, we compute $\texttt{shrinkf}(f_1^{(i)}, q_2, e_2')$ for each $i \in \{1, \ldots, t_1\}$ and concatenate the results to obtain $\big(f_2^{(1)}(x), \ldots, f_2^{(t_2)}(x)\big)$ so that each component of each $\texttt{shrinkf}(f_1^{(i)}, q_2, e_2')$ occurs as some $f_2^{(i)}(x)$ (multiplicities discarded). We continue in this manner until we obtain $\big(f_s^{(1)}(x), \ldots, f_s^{(t_s)}(x)\big)$, the values of $\texttt{shrinkf}(f_{s-1}^{(i)}, q_s, e_s')$ concatenated. The revised algorithm now reads the same as before except, in the case that $f_s(1) = 0$, we do not check if $\Phi_d(x) | f(x)$ by a direct computation. Instead we check if every $f_s^{(i)}(1) = 0$. We claim that $\Phi_d(x) | f(x)$ if and only if $f_s^{(i)}(1) = 0$ for every $i \in \{1, 2, \ldots, t_s\}$. The definition of `shrinkf` and the converse of Corollary 1 imply that if every $f_s^{(i)}(1) = 0$, then each $f_{s-1}^{(i)}(x)$ is divisible by $\Phi_{q_s^{e_s'}}(x)$. Now, the converse of Theorem 1 implies that each $f_{s-2}^{(i)}(x)$ is divisible by $\Phi_{q_{s-1}^{e'_{s-1}} q_s^{e_s'}}(x)$, and continued applications of the converse of Theorem 1 imply $\Phi_d(x) | f(x)$. The argument to see that $\Phi_d(x) | f(x)$ implies every $f_s^{(i)}(1) = 0$ is similar but in reverse. Thus, if (6) holds, we simply perform the algorithm as described in Section 2 replacing each check to see if $\Phi_d(x) | f(x)$ by checking instead if $f_s^{(i)}(1) = 0$ for every $i \in \{1, 2, \ldots, t_s\}$. (As noted there some savings is gained by using information obtained from previous calls to `shrinkf`; in particular, for any given d as above, the values of $f_j^{(i)}(x)$ with $j < s$ will be known prior to considering d.)

An estimate for the running time of the algorithm as revised here can be seen to depend heavily on the number of different $f_j^{(i)}(x)$ that need to be computed. If the input polynomial is $f(x) = \sum_{j=1}^{N} a_j x^{d_j}$ where the a_j are nonnegative integers and the d_j are integers satisfying $0 \le d_1 < d_2 < \cdots < d_N$, then a deterministic upper bound for the running time is

$$\ll \exp\left((2 + o(1))\sqrt{N}\big(\sqrt{\log N} + \log\log d_N\big)\right) \times \log \max_{1 \le j \le N} \{|a_j| + 1\}$$

as N tends to infinity. This estimate is hampered mainly by a presumably poor estimate for the number of elements of S skipped over in the approach described in Section 2. In any case, with N fixed, one sees that the running time depends only logarithmically on the degree of $f(x)$.

We note that H. W. Lenstra [2, Proposition 3.5] has described an algorithm which takes a given lacunary polynomial and determines all cyclotomic factors, with their multiplicities, having degree less than some prescribed amount. It would be of interest to find an efficient algorithm for determining all the cyclotomic divisors of a lacunary polynomial $f(x)$. We would also be interested in good upper bounds for the number of m considered in the algorithms described here (excluding the m we skipped over by considerations of an exact divisor d of m).

References

1. J. H. Conway and A. J. Jones, *Trignometric diophantine equations (On vanishing sums of roots of unity)*, Acta Arith. **30** (1976), 229–240.
2. H. W. Lenstra, *Finding small degree factors of lacunary polynomials*, preprint.
3. H. B. Mann, *On linear relations between roots of unity*, Mathematika **12** (1965), 107–117.