Fast Rational Function Reconstruction *

Sara Khodadad
School of Computing Science,
Simon Fraser University,
Burnaby, B.C. V5A 1S6, CANADA.
skhodada@cecm.sfu.ca.

Michael Monagan
Department of Mathematics,
Simon Fraser University,
Burnaby, B.C. V5A 1S6, CANADA.
mmonagan@cecm.sfu.ca.

ABSTRACT

Let F be a field and let f and g be polynomials in $F[t]$ satisfying $\deg f > \deg g$. Recall that on input of f and g the extended Euclidean algorithm computes a sequence of polynomials (s_i, t_i, r_i) satisfying $s_i f + t_i g = r_i$. Thus for i with $\gcd(t_i, f) = 1$, we obtain rational functions $r_i/t_i \in F(t)$ satisfying $r_i/t_i \equiv g \pmod f$.

In this paper we modify the fast extended Euclidean algorithm to compute the smallest r_i/t_i, that is, an r_i/t_i minimizing $\deg r_i + \deg t_i$. This means that in an output sensitive modular algorithm when we are recovering rational functions in $F(t)$ from their images modulo $f(t)$ where $f(t)$ is increasing in degree, we can recover them as soon as the degree of f is large enough and we can do this fast.

We have implemented our modified fast Euclidean algorithm for $F = \mathbb{Z}_p$, p a word sized prime, in Java. Our fast algorithm beats the ordinary Euclidean algorithm around degree 200. This has application to polynomial gcd computation and linear algebra over $\mathbb{Z}_p(t)$.

1. INTRODUCTION

Rational number reconstruction, originally developed by Paul Wang in [15], (see [2] or [14] for an accessible reference), has found many applications in computer algebra. It enables us to design efficient *modular algorithms* for computing with polynomials, vectors and matrices over $\mathbb{Q}$. Such algorithms first solve a problem modulo a sufficiently large integer m which is usually a product of primes or a power of a prime. Then they apply rational reconstruction to recover the rational numbers in the solution from their images modulo m. The same basic strategy can also be used to recover fractions in $F(t)$ from their image modulo a polynomial $f(t) \in F[t]$ where F is a field. Some applications where rational reconstruction has been used include polynomial gcd computation over $\mathbb{Q}(\alpha)$, solving linear systems over $\mathbb{Q}$ and Gröbner basis computation over $\mathbb{Q}$.

*Supported by NSERC of Canada and the MITACS NCE of Canada

A key advantage of rational reconstruction is that it enables us to make modular algorithms "output sensitive", that is, the size of the modulus m needed, and hence overall efficiency, depends on the size of the actual rationals in the output and not on bounds for their size which might be much larger. For example, consider the problem of computing the monic gcd g of two non-zero univariate polynomials f_1 and f_2 in $L[x]$ where $L = \mathbb{Q}(\alpha)$ is a number field with minimal polynomial $m(z)$. Let $\mathrm{lc}(f) \in L$ denote the leading coefficient of f. Let $\mathrm{den}(f) \in \mathbb{Z}$ denote the lcm of the denominators in the rational coefficients of f and let $\hat{f} = \mathrm{den}(f)f$. Thus $\mathrm{lc}(g) = 1$ and $\hat{g}$ has no fractions in the coefficients. A modular algorithm for the gcd problem in $L[x]$ for the case $m(z) \in \mathbb{Z}[z]$. was first given by Langemyr and MacCallum in [5] Their algorithm computes $g_1, g_2, ...$ the gcd of f_1 and f_2 modulo a sequence of primes $p_1, p_2, ...$ using the Euclidean algorithm and applies the Chinese remainder theorem to reconstruct Dg from the images $Dg_i \bmod p_i$ where D is a known multiple of $\mathrm{den}(g)$. They use

$$D = \Delta \times \mathrm{GCD}(\mathrm{den}((\mathrm{lc}\hat{f}_1)^{-1}), \mathrm{den}((\mathrm{lc}\hat{f}_2)^{-1})))$$

where $\Delta = \mathrm{res}_z(m(z), m'(z))$ is the discrimant of the number field. But g might be the polynomial $x+2\alpha+5/3$ and the inputs f_1 and f_2 and $m(z)$ might have large integers in their coefficients. In such cases, the integer coefficients of Dg can be much larger than those of $\mathrm{den}(g)g = 3g = 3x + 6\alpha + 5$.

In [3], Encarnacion modified Langemyr and MacCallum's algorithm to work for $m(z) \in Q[z]$ and to use rational number reconstruction to make it output sensitive. For $i = 1, 2, ...$ Encarnacion computes $c_i = g \bmod m_i$ from $g_1, g_2, ..., g_i$ where $m_i = p_1 \times ... \times p_i$ and applies rational reconstruction to the coefficients of c_i to obtain h_i. When the output of rational reconstruction succeeds and does not change from one prime to the next, then $h_i = g$ with high probability. One verifies this by trial division, that is, by checking $h_i | f_1$ and $h_i | f_2$. Thus for $g = x+2\alpha+5/3$, only two primes would be needed regardless of how large $f_1(x), f_2(x)$ and $m(z)$ are.

Wang's Algorithm

Let $n/d \in \mathbb{Q}$ with $d > 0$ and $\gcd(n, d) = 1$. Let $m \in \mathbb{Z}$ with $m > 0$ and $\gcd(m, d) = 1$. Suppose we have computed $u = n/d \bmod m$ and we want to recover the rational n/d. Recall that extended Euclidean algorithm (EEA), on input of m and u with $m > u \geq 0$, computes a sequence of triples $(s_i, t_i, r_i) \in \mathbb{Z}^3$ for $i = 0, 1, \ldots, l, l+1$ satisfying $r_{l+1} = 0$ and $s_i m + t_i u = r_i$. It does this by initializing $(r_0, s_0, t_0) = (m, 1, 0)$ and $(r_1, s_1, t_1) = (u, 0, 1)$ and computing

$$(r_{i+1}, s_{i+1}, t_{i+1}) = (r_{i-1} - q_i r_i, s_{i-1} - q_i s_i, t_{i-1} - q_i t_i)$$

for $i = 1, 2, \ldots, l$ where q_i is the quotient of r_{i-1} divided r_i. Observe that $s_i m + t_i u = r_i$ implies $r_i/t_i \equiv u \pmod m$ for all i with $\gcd(m, t_i) = 1$. In [15], Wang observed that if $m > 2|n|d$ then the rational $n/d = r_i/t_i$ for some $0 \le i \le l+1$. In fact, it is the r_i/t_i satisfying $r_{i-1} > |n| \ge r_i$, that is, we just need to compute up to the first remainder less than or equal to $|n|$.

One way to use Wang's observations to recover the rational number n/d in the output from its image u modulo m is as follows. First bound the size of n and d, that is, compute $N \ge |n|$ and $D \ge d$. Then solve the problem modulo a sequence of primes $p_1, p_2, \ldots$ satisfying $m > 2ND$ where $m = p_1 \times p_2 \times \ldots$. Then run the Euclidean algorithm until $r_{i-1} > N \ge r_i$, and output r_i/t_i after checking that $\gcd(t_i, m) = 1$.

However, as remarked earlier, the bounds are often much too big. To make a modular algorithm output sensitive we let m increase in size and *periodically* apply rational reconstruction as follows. Given the image u of the rational n/d modulo m, Wang computes $N = D = \lfloor\sqrt{m/2}\rfloor$ and runs the Euclidean algorithm with input $m > u$ stopping when $r_{i-1} > N \ge r_i$. One then checks if $|t_i| \le D$ and $\gcd(t_i, m) = 1$. If yes then we output r_i/t_i else rational reconstruction "fails". Thus Wang's algorithm succeeds in reconstructing n/d when m becomes bigger than $2\max(n^2, d^2)$.

If one uses the ordinary Euclidean algorithm, the complexity of Wang's algorithm is $O(\log^2 m)$. In 2002 Pan and Wang in [11] modified the fast Euclidean algorithm of Schönhage [12] to solve the rational number reconstruction problem in time $O(\mathsf{M}(k)\log k)$ where $k = \log m$ is the length of the modulus m and $\mathsf{M}(k)$ is the cost of multiplying integers of length k. The authors did not implement their algorithm and remarked during their presentation at ISSAC 2002 that the algorithm might not be practical. In 2005 Lichtblau in [6] implemented a variation on the fast Euclidean algorithm for rational number reconstruction for Mathematica and found that it is practical. In fact, Steel (see [13]) had already implemented fast rational number reconstruction in Magma version 2.8 in 2000. Steel found that the fast Euclidean algorithm beat his implementation of Lehmer's algorithm (an improvement of a constant factor on the ordinary Euclidean algorithm) for integers of length around 50,000 bits.

Maximal Quotient Rational Reconstruction

There is an inefficiency in Wang's approach because of the choice of $N = D = \lfloor\sqrt{m/2}\rfloor$. This choice means we are using half of the bits of m to recover the numerator and half for the denominator. To recover n/d, we require $m > 2|n|d$ but this choice for N and D means the modulus $m > 2\max(n^2, d^2)$. This choice is efficient if the numerator n and denominator d are of the same length. But if $|n| \gg d$ or $|n| \ll d$, it requires m to be up to twice as long as is necessary. This inefficiency was noted by Monagan in [9]. In particular, for gcd problems in $L[x]$, Monagan has observed that the denominators in g are often much smaller than numerators.

Monagan in [9] observed that if $m \gg 2|n|d$ then with high probability there will be only one small rational r_i/t_i in the Euclidean algorithm, namely n/d. In fact, if m is just a few bits longer than $2|n|d\log_2 m$, the smallest rational will be n/d with high probability. Thus another way to solve the rational reconstruction problem is to simply select and output the smallest r_i/t_i. How do we do this without explicitly multiplying $r_i \times t_i$? Monagan observed that if the size of the rational r_i/t_i is small compared with m, that is, $|r_i t_i| \ll m$ then $q_i = \lfloor r_{i-1}/r_i \rfloor$ is necessarily large, indeed q_i satisfies $\frac{m}{3} < q_i r_i |t_i| \le m$. Hence, it is sufficient to select the rational r_i/t_i corresponding to the largest quotient q_i. Moreover, since the quotients are available and they are mostly very small integers, this selection is efficient.

In this way, it does not really matter whether n is much longer or much shorter than d, for as soon as m is a few bits longer than $2|n|d\log_2 m$, we can select n/d from the r_i/t_i with high probability. If m is a product of primes and one is using the Chinese remainder theorem, one saves up to half the number of primes. Thus in an application where the size of the numerators might be much larger or smaller than the size of the denominators, Monagan's algorithm is preferred.

Monagan's algorithm, like Wang's algorithm, is also a simple modification of the extended Euclidean algorithm, and thus also has complexity $O(\log^2 m)$ if the ordinary extended Euclidean algorithm is used. Just as Pan and X. Wang modified the fast Euclidean algorithm to accelerate Wang's algorithm, can Monagan's algorithm also be accelerated? In this paper we answer this question in the affirmative. We show how to modify the fast Euclidean algorithm to output the smallest rational r_i/t_i without increasing the asymptotic time complexity of the fast Euclidean algorithm. Rather than modifying the fast Euclidean algorithm for $\mathbb{Z}$, we modify the fast Euclidean algorithm for $\mathbb{Z}_p[t]$ where p is a prime to recover the rational function for $\mathbb{Z}_p(t)$ of least degree.

We call our algorithm FMQRFR for fast maximal quotient rational function reconstruction. We have implemented it in Java. In comparing it to an implementation using the ordinary extended Euclidean algorithm for $\mathbb{Z}_p[t]$ we found that the fast Euclidean algorithm beats the ordinary extended Euclidean algorithm at around degree 200. In order to achieve such a result, one must implement fast multiplication in $\mathbb{Z}_p[t]$ carefully. For this we have implemented an "in-place" version of Karatsuba's algorithm (see Maeder [7]) so that fast multiplication in $\mathbb{Z}_p[t]$ already beats classical multiplication at degree 50.

We also show how to accelerate Wang's algorithm for $\mathbb{Z}_p[t]$ using the fast Euclidean algorithm. Given $m(t)$, the cost of FMQRFR is no more than twice as expensive as the accelerated Wang's algorithm. If one uses Wang's algorithm, one runs the (fast) Euclidean algorithm half way, that is, stopping as soon as the degree of the remainder is below $(\deg m)/2$. In our algorithm, one cannot stop in the middle in general. One must run through to the end computing all quotients to determine the largest. Hence the factor of 2.

Our paper is organized as follows. In section 2 we describe the maximal quotient rational reconstruction algorithm for $F[x]$. In section 3 we describe the fast extended Euclidean algorithm (FEEA) for $F[x]$. Our presentation of the FEEA follows the presentation given by von zur Gathen and Gerhard in [14]. We give timings for our implementation of the FEEA for $F = \mathbb{Z}_p$ where p is a word size prime, comparing it with the ordinary extended Euclidean algorithm. In section 4 we show how to modify the FEEA to compute the smallest rational function r_i/t_i. The reason this is possible is that the FEEA computes all the quotients q_i. In section 4 we also show how to modify the FEEA to compute the rational function r_i/t_i satisfying $\deg r_{i-1} > (\deg f)/2 \ge \deg r_i$ so that we can compare the efficiency of our algorithm with Wang's algorithm.

2. MAXIMAL QUOTIENT RATIONAL FUNCTION RECONSTRUCTION

Let F be a field. A rational function $n/d \in F(x)$ is said to be in *canonical form* if $\mathrm{lc}(d) = 1$ and $\gcd(n, d) = 1$. Let $f, g \in F[x]$ with $\deg f > \deg g$. Let r_i and t_i be the elements of the ith row of the Extended Euclidean Algorithm (EEA) with inputs f and g. Then any rational function n/d with $n = r_i/\mathrm{lc}(t_i)$ and $d = t_i/\mathrm{lc}(t_i)$ satisfies $n/d \equiv g \mod f$, provided that $\gcd(f, t_i) = 1$. Moreover, if n/d is a canonical form solution to $n/d \equiv g \mod f$ satisfying $\deg n + \deg d < \deg f$, then there exists some row j in the EEA for inputs f and g such that $n = r_j/\mathrm{lc}(t_j)$ and $d = t_j/\mathrm{lc}(t_j)$. Thus the EEA with inputs f and g generates all rational functions n/d (up to scalar multiples in F) satisfying $n/d \equiv g \mod f$, $\gcd(f, d) = 1$ and $\deg n + \deg d < \deg f$. Refer to [14, Lemma 5.15] for the proof.

If degree bounds $N \geq \deg n$ and $D \geq \deg d$ satisfying $N + D < \deg f$ are known, then the rational function n/d is uniquely determined by running the EEA on inputs f and g. But we do not always know the values of N and D in advance. In this section we will present an efficient algorithm that with high probability finds the correct solution for $\deg f > \deg n + \deg d + 1$. The following example helps us in the design of the algorithm.

EXAMPLE 2.1. *Consider* $f = \prod_{i=5}^{12}(x - i)$ *and* $g = 10x^7 + x^6 + 2x^5 + 10x^4 + 12x^3 + 7x^2 + 12x + 8$ *in* $\mathbb{Z}_{13}[x]$. *The Extended Euclidean Algorithm with inputs f and g yields the following table.*

i	$\deg r_i$	$\deg t_i$	$\deg r_i + \deg t_i$	$\deg q_i$
1	7	0	7	1
2	6	1	7	1
3	5	2	7	1
4	2	3	5	3
5	1	6	7	1
6	0	7	7	1

The data in the table suggest that we simply return the rational function r_i/t_i where $\deg r_i + \deg t_i$ is minimal. As illustrated in the table, r_4/t_4 has minimal total degree of 5. Note that it also corresponds to the quotient of maximal degree q_4. The reason for this is easily explained by the following lemma.

LEMMA 2.2. *Let F be a field and $f, g \in F[x]$. In the EEA for f and g we have*

$$\deg r_i + \deg t_i + \deg q_i = \deg f$$

for $1 \leq i \leq l$ where l is the total number of division steps in the EEA for inputs f and g.

PROOF. We know $\deg t_i = \deg f - \deg r_{i-1}$, thus

$$\deg r_i + \deg t_i + \deg q_i =$$
$$\deg r_i + (\deg f - \deg r_{i-1}) + \deg r_{i-1} - \deg r_i = \deg f.$$

□

The following lemma states that when $\deg f$ is large enough then there would only be one pair of (r_j, t_j) such that $\deg r_j + \deg t_j$ is minimal.

LEMMA 2.3. *Let F be a field, and $n, d \in F[x]$ with $\mathrm{lc}(d) = 1$ and $\gcd(n, d) = 1$. Let f, g be two polynomials in $F[x]$ satisfying $\gcd(f, d) = 1$ and $g = n/d \mod f$. Let q_j denote a quotient with maximal degree in the Extended Euclidean Algorithm with inputs f and g. If $\deg f > 2(\deg n + \deg d)$ then q_j is unique, $n = r_j$ and $d = t_j$.*

PROOF. Since $\deg f > \deg n + \deg d$, then in the Extended Euclidean Algorithm with inputs f and g there exists a unique index j such that $r_j/t_j = n/d$. Thus according to Lemma 2.2 we have $\deg q_j > 1/2 \deg f$. On the other hand, we know $\sum_{i=1}^{l} \deg q_i = \deg f - \deg r_l \leq \deg f$ which implies that q_j is the quotient with maximal degree or $\deg r_j + \deg t_j$ is minimal. Therefore $n = r_j$ and $d = t_j$ with j being the index for which $\deg r_j + \deg t_j$ is minimal. □

Maximal Quotient RFR Algorithm (MQRFR)
Input: $f, g \in \mathbb{Z}_p[x]$ with $\deg f > \deg g$, and $T \in \mathbb{N}$
Output: Either $n, d \in \mathbb{Z}_p[x]$ satisfying $n/d \equiv g \mod f$, $\mathrm{lc}(d) = 1$, $\gcd(n, d) = 1$, and $\deg n + \deg d + T < \deg f$, or FAIL implying no solution exists

1. if $g = 0$ then if $\deg f \geq T$ then return $(0, 1)$ else return FAIL
2. $(r_0, r_1) = (f, g)$
   $(t_0, t_1) = (0, 1)$
   $(n, d) = (r_1, t_1)$
3. while $r_1 \neq 0$ do
   - if $\deg n + \deg d > \deg r_1 + \deg t_1$ then
     - $(n, d) = (r_1, t_1)$
   - $q = r_0 \operatorname{quo} r_1$
   - $(r_0, r_1) = (r_1, r_0 - qr_1)$
   - $(t_0, t_1) = (t_1, t_0 - qt_1)$
4. if $\deg n + \deg d + T \geq \deg f$ or $\gcd(n, d) \neq 1$ then
   return FAIL
5. return $(n/\mathrm{lc}(d), d/\mathrm{lc}(d))$

If we let $m = \deg f$, then assuming classical algorithms are used for multiplication, division and GCD computation in $\mathbb{Z}_p[x]$, step 3 and step 4 take $O(m^2)$ operations in $\mathbb{Z}_p$. Thus MQRFR is of time complexity $O(m^2)$.

3. THE FAST EUCLIDEAN ALGORITHM

In 1971 Schönhage in [12] presented a fast integer GCD algorithm with time complexity $O(n \log^2 n \log \log n)$. An asymptotically fast rational number reconstruction algorithm based on Schönhage's algorithm was presented by Pan and Wang in [11]. Before that Allan Steel had implemented in Magma a fast rational number reconstruction algorithm based on the half-gcd algorithm presented in Montgomery's PhD thesis [10] for polynomials in $F[x]$. Maple v. 10, Mathematica v. 4.0 and Magma v. 2.10 all have fast integer multiplication and division implemented, but Magma is the only computer algebra system having implemented the fast integer GCD algorithm.

Assuming a multiplication algorithm of time complexity $O(n \log^a n)$ is available for polynomials of size n in $F[x]$, in 1973 Moenck in [8] adapted Shoñhage's algorithm into an $O(n \log^{a+1} n)$ algorithm for polynomial GCD computation in $F[x]$. In 1980 Brent, Gustavson, and Yun in [1] gave two speedups for Moenck's algorithm. They also pointed out (but did not prove) a generalization of Moenck's algorithm. Later in 1992, Montgomery in his PhD thesis [10] independently stated and proved a similar generalization of Moenck's algorithm with some of the same speedups.

In this section we describe the Fast Euclidean Algorithm presented by von zur Gathen and Gerhard in [14] and then in the next section we show how it can be modified to be used in the Maximal Quotient Rational Function Reconstruction Algorithm. Thus we follow [14] in our assumptions and definitions.

Let F be a field and $r_0, r_1 \in F[x]$ with $\deg r_0 \geq \deg r_1$. Let

$$\begin{aligned} \rho_{i+1} r_{i+1} &= r_{i-1} - q_i r_i, \\ \rho_{i+1} s_{i+1} &= s_{i-1} - q_i s_i, \\ \rho_{i+1} t_{i+1} &= t_{i-1} - q_i t_i, \end{aligned}$$

for $1 \leq i \leq l$, be the results of the Extended Euclidean Algorithm for r_0, r_1, where $s_0 = t_1 = 1$, $s_1 = t_0 = 0$ and $r_{l+1} = 0$. We let ρ_i to denote the leading coefficient of the remainders, that is $\operatorname{lc}(r_i) = 1$. Let $R_i = Q_i \ldots Q_1 R_0$, for $1 \leq i \leq l$, where

$$Q_i = \begin{bmatrix} 0 & 1 \\ 1/\rho_{i+1} & -q_i/\rho_{i+1} \end{bmatrix}, \quad R_0 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

in $F[x]^{2\times 2}$. Then it can be easily proved by induction on i that

$$R_i = \begin{bmatrix} s_i & t_i \\ s_{i+1} & t_{i+1} \end{bmatrix}.$$

This matrix is of great importance in the design of the Fast Extended Euclidean Algorithm.

Let $f = f_n x^n + f_{n-1} x^{n-1} + \ldots + f_0 \in F[x]$ and $f_n \neq 0$. We define the truncated polynomial

$$f \upharpoonright k = f \operatorname{quo} x^{n-k} = f_n x^k + f_{n-1} x^{k-1} + \ldots + f_{n-k},$$

for $k \in \mathbb{Z}$. The polynomial $f \upharpoonright k$ is of degree k for $k \geq 0$ and its coefficients are the $k+1$ highest coefficients of f. The pairs (f, g) and (f^*, g^*) *coincide up to* k if

$$\begin{aligned} f \upharpoonright k &= f^* \upharpoonright k, \\ g \upharpoonright (k - (\deg f - \deg g)) &= g^* \upharpoonright (k - (\deg f^* - \deg g^*)), \end{aligned}$$

where $f, g, f^*, g^* \in F[x] \backslash \{0\}$, $\deg f \geq \deg g$, $\deg f^* \geq \deg g^*$ and $k \in \mathbb{Z}$. Following [14], we define the positive integer number $\eta_{f,g}(k)$ for any $k \in \mathbb{N}$ and $f, g \in F[x]$ by

$$\eta_{f,g}(k) = \max_{0 \leq j \leq l} \{ j : \sum_{i=1}^{j} m_i \leq k \},$$

where l denotes the number of division steps in the Euclidean algorithm with inputs f and g. The following lemma implies that the first $\eta_{f,g}(k)$ results of the Euclidean Algorithm only depend on the top part of the inputs, which is the basic idea leading to a fast GCD algorithm.

LEMMA 3.1. *[14, Lemma 11.3] Let $k \in \mathbb{N}$, $h = \eta_{r_0,r_1}(k)$ and $h^* = \eta_{r_0^*,r_1^*}(k)$, with r_0, r_1, r_0^*, r_1^* monic polynomials in $F[x]$. If (r_0, r_1) and (r_0^*, r_1^*) coincide up to $2k$ and $k \geq \deg r_0 - \deg r_1$, then*

1. *$h = h^*$,*
2. *$q_i = q_i^*$ for $1 \leq i \leq h$,*
3. *$\rho_i = \rho_i^*$ for $2 \leq i \leq h$,*

where $q_i, q_i^* \in F[x]$ and $\rho_i, \rho_i^* \in F$ are defined by

$$\begin{aligned} r_{i-1} &= q_i r_i + \rho_{i+1} r_{i+1} \quad (1 \leq i \leq l), \quad r_{l+1} = 0, \\ r_{i-1}^* &= q_i^* r_i^* + \rho_{i+1}^* r_{i+1}^* \quad (1 \leq i \leq l^*), \quad r_{l^*+1}^* = 0. \end{aligned}$$

Refer to [14] for a detailed proof of this lemma. To improve the efficiency of the EEA a divide-and-conquer algorithm, called *Fast Extended Euclidean Algorithm*, is designed based on the above lemma. Von zur Gathen and Gerhard in [14, Ch. 11] present Shönhage's Fast Extended Euclidean Algorithm for polynomials in $F[x]$, however, the algorithm presented in the book needs some minor corrections. We asked the authors to send us the correct version of the algorithm which is described by the following algorithm, however, we have modified it by removing some unnecessary outputs in the following.

Fast Extended Euclidean Algorithm (FEEA)
Input: r_0 and r_1 two *monic* polynomials in $F[x]$ with $n_0 = \deg r_0 > n_1 = \deg r_1 \geq 0$ and $k \in \mathbb{N}$ with $n_0/2 \leq k \leq n_0$
Output: $h = \eta_{r_0,r_1}(k) \in \mathbb{N}$, $\rho_{h+1} \in F$, $R_h = \begin{bmatrix} s_h & t_h \\ s_{h+1} & t_{h+1} \end{bmatrix}$

1. if $r_1 = 0$ or $k < n_0 - n_1$ then
   return 0, 1, $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$
   else if $n_0 <$ *cutoff* then
   return EEA(r_0, r_1, k)
2. $k_1 = \lfloor k/2 \rfloor$
   $r_0^* = r_0 \upharpoonright 2k_1$, $r_1^* = r_1 \upharpoonright (2k_1 - (n_0 - n_1))$
   $j-1$, ρ_j^*, R_{j-1}^*= FEEA(r_0^*, r_1^*, k_1)
3. compute $\rho_j, R_{j-1}, r_{j-1}, r_j$ and $n_j = \deg r_j$
4. if $r_j = 0$ or $k < n_0 - n_j$ then
   return $j-1$, ρ_j, R_{j-1}
5. compute $q_j, \rho_{j+1}, r_{j+1}, n_{j+1} = \deg r_{j+1}$ and
   $R_j = \begin{bmatrix} 0 & 1 \\ 1/\rho_{j+1} & -q_j/\rho_{j+1} \end{bmatrix} R_{j-1}$
6. $k_2 = k - (n_0 - n_j)$
   $r_j^* = r_j \upharpoonright 2k_2$, $r_{j+1}^* = r_{j+1} \upharpoonright (2k_2 - (n_j - n_{j+1}))$
   $h-j$, ρ_{h+1}^*, S^*= FEEA(r_j^*, r_{j+1}^*, k_2)
7. compute ρ_{h+1}, S, r_h and r_{h+1}
8. return h, ρ_{h+1}, SR_j

As illustrated above, besides the two monic polynomials r_0 and r_1, the algorithm gets a third input $k \in \mathbb{N}$. A sequence of quotients is recursively computed as in the Extended Euclidean Algorithm, until the some of the degrees of the quotients is greater than k. That is, if $h = \eta_{r_0,r_1}(k)$ denotes the last computed quotient, then we will have

$$\sum_{i=1}^{h} \deg q_i \leq k < \sum_{i=1}^{h+1} \deg q_i.$$

The FEEA divides the problem into two subproblems of almost the same size, i.e., the sum of the degrees of the quotients computed in each recursive call is less than or equal to $k/2$. Note that in this algorithm all elements of the EEA are computed except the remainders. However, having s_h and t_h as the entries of the second row of the matrix R_h one can easily compute a single remainder r_h by writing

$r_h = s_h r_0 + t_h r_1$. It is not hard to see that $r_h = \gcd(r_0, r_1)$, if we set $k = \deg r_0$.

According to Lemma 3.1, ρ_j is not necessarily equal to ρ_j^*, and thus R_{j-1} and R_{j-1}^* are not equal either. Therefore we use the following relations

$$\begin{bmatrix} r_{j-1} \\ \tilde{r}_j \end{bmatrix} = R_{j-1}^* \begin{bmatrix} r_0 \\ r_1 \end{bmatrix}, \quad R_{j-1} = \begin{bmatrix} 1 & 0 \\ 0 & 1/\mathrm{lc}(\tilde{r}_j) \end{bmatrix} R_{j-1}^*,$$
$$\rho_j = \rho_j^* \mathrm{lc}(\tilde{r}_j), \quad r_j = \tilde{r}_j / \mathrm{lc}(\tilde{r}_j),$$

in step 3 to compute ρ_j, R_{j-1}, r_{j-1} and r_j. Similar computations are performed in step 7 to compute ρ_{h+1}, S, r_h and r_{h+1}. The algorithm has a time complexity of $O(\mathsf{M}(k) \log k)$, where $\mathsf{M}(k)$ denotes the number of field operation required to multiply two univariate polynomials of degree k. Refer to [4, p. 27] for a detailed cost analysis of the algorithm and a detailed proof of correctness of the algorithm.

We have implemented the FEEA for polynomials in $F[x] = \mathbb{Z}_p[x]$ in Java. We used Karatsuba's algorithm for univariate polynomial multiplication in our implementation which is of time complexity $O(n^{\log_2 3})$ for polynomials of degree n. The algorithm is not effective in practice for polynomials of low degree. We use the classical multiplication method for polynomials of degree less than 50 and switch to Karatsuba's when the input polynomials have a degree greater than 50. The following table includes the timings we gathered for our implementation of the Classical and Karatsuba's multiplication algorithms over $\mathbb{Z}_p[x]$, where p is a 15 bit prime. As illustrated below, the timings of Karatsuba's algorithm increase by a factor close to 3 as the degree doubles which confirms that our implementation is of time complexity $O(n^{\log_2 3})$.

n	Karatsuba(ms)	Classical(ms)
128	0.34	0.38
256	0.98	1.40
512	2.93	5.40
1024	8.93	21.62
2048	26.48	84.43
4096	79.78	345.67
8192	245.04	1375.42

It turns out that in practice the EEA performs better than the FEEA as well for polynomials of low degree. Our implementation of the FEEA beats the EEA when $\deg r_0 = 200$. Thus we use the EEA for dividends of degree below 200 in step 1 of the FEEA. The following figure illustrates the timings (in ms) of the FEEA on two random polynomials of degree 10000 for different cutoff degrees.

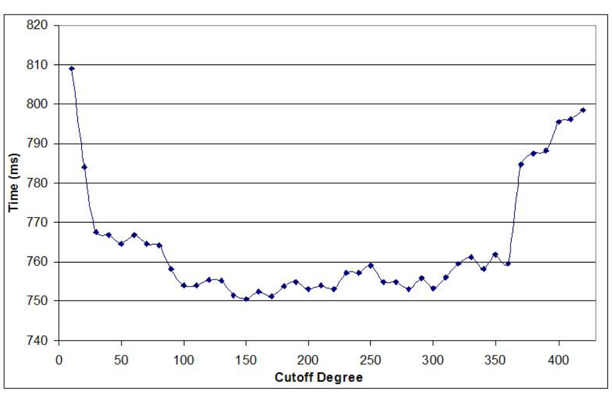

Our Java implementation of the EEA accepts 3 inputs and returns the same outputs as the FEEA. The following table includes our timings for the EEA and the FEEA on random polynomials of degree n. It shows that we see a significant speedup by $n = 1000$.

n	EEA(ms)	FEEA(ms)	r_1	r_2
1000	373.80	295.63	0.00052	1.26
2000	1427.18	942.83	0.00050	1.51
4000	5602.18	2972.08	0.00049	1.88
8000	22295.47	9588.76	0.00048	2.33
16000	88766.90	31278.50	0.00049	2.84
32000	354085.71	99273.77	0.00048	3.54

$r_1 = \mathrm{FEEA}/(n^{\log_2 3} \log n)$, $r_2 = \mathrm{EEA}/\mathrm{FEEA}$

4. MQRFR USING FEEA

To make the MQRFR algorithm more efficient we use the FEEA instead of the EEA. As pointed out before, the FEEA does not compute the intermediate remainders, but it does compute all the quotients. Also s_i and t_i are available as the entries of the first row of R_i. Thus according to lemma 2.2 instead of selecting r_i and t_i such that $\deg r_i + \deg t_i$ is minimal, we can return q_i the quotient with maximal degree along with corresponding values of s_i and t_i. The remainder r_i is then obtained from s_i and t_i using two long multiplications ($r_i = s_i f + t_i g$). The following algorithm presents the FEEA modified to return the quotient of maximal degree.

Modified FEEA (MFEEA)
Input: r_0 and r_1 two *monic* polynomials in $F[x]$ with $n_0 = \deg r_0 > n_1 = \deg r_1 \geq 0$ and $k \in \mathbb{N}$ with $n_0/2 \leq k \leq n_0$
Output: $h = \eta_{r_0,r_1}(k) \in \mathbb{N}$, $\rho_{h+1} \in F$, $R_h = \begin{bmatrix} s_h & t_h \\ s_{h+1} & t_{h+1} \end{bmatrix}$, $q_{\max}$, $s_{\max}$, $t_{\max}$

1. if $r_1 = 0$ or $k < n_0 - n_1$ then
   return $0,\ 1,\ \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},\ 1,\ 1,\ 0$
   else if $n_0 <$ *cutoff* then
   return $\mathrm{EEA}(r_0,\ r_1,\ k)$
2. $k_1 = \lfloor k/2 \rfloor$
   $r_0^* = r_0 \upharpoonright 2k_1, r_1^* = r_1 \upharpoonright (2k_1 - (n_0 - n_1))$
   $j-1,\ \rho_j^*,\ R_{j-1}^*,\ q_{\max},\ s_{\max},\ t_{\max} = \mathrm{MFEEA}(r_0^*,\ r_1^*, k_1)$
3. compute $\rho_j, R_{j-1}, r_{j-1}, r_j$ and $n_j = \deg r_j$
4. if $r_j = 0$ or $k < n_0 - n_j$ then
   return $j-1,\ \rho_j,\ R_{j-1},\ q_{\max},\ s_{\max},\ t_{\max}$
5. compute $q_j, \rho_{j+1}, r_{j+1}, n_{j+1} = \deg r_{j+1}$ and
   $R_j = \begin{bmatrix} 0 & 1 \\ 1/\rho_{j+1} & -q_j/\rho_{j+1} \end{bmatrix} R_{j-1}$
   if $\deg q_j > \deg q_{\max}$ then
   $q_{\max}, s_{\max}, t_{\max} = q_j, R_j[1,1], R_j[1,2]$
6. $k_2 = k - (n_0 - n_j)$
   $r_j^* = r_j \upharpoonright 2k_2,\ r_{j+1}^* = r_{j+1} \upharpoonright (2k_2 - (n_j - n_{j+1}))$
   $h-j,\ \rho_{h+1}^*,\ S^*,\ q_{\max}^*,\ s_{\max}^*,\ t_{\max}^* = \mathrm{MFEEA}(r_j^*, r_{j+1}^*, k_2)$
   if $\deg q_{\max}^* > \deg q_{\max}$ then
   $q_{\max} = q_{\max}^*$
   $\begin{bmatrix} s_{\max} & t_{\max} \end{bmatrix} = \begin{bmatrix} s_{\max}^* & t_{\max}^* \end{bmatrix} R_j$
7. compute ρ_{h+1}, S, r_h and r_{h+1}
8. return $h,\ \rho_{h+1},\ SR_j,\ q_{\max},\ s_{\max},\ t_{\max}$

As illustrated above the only modification we have made to the FEEA is to return three more outputs, i.e., $q_{\max}$, $s_{\max}$, $t_{\max}$. Thus assuming the FEEA works correctly, we require to prove that $q_{\max}$ is the quotient with maximal degree and $s_{\max}$ and $t_{\max}$ have the same index as $q_{\max}$ in the Euclidean Algorithm with inputs r_0 and r_1.

We see by induction on k that the results of the recursive call in step 2 are correct, that is, $q_{\max}$ represents the quotient with maximal degree in $\{q_1, \ldots, q_{j-1}\}$ and $s_{\max}$ and $t_{\max}$ are in the same row with $q_{\max}$. In step 4 the correct result is returned, since no other quotient has been computed. We have

$$R_j = \begin{bmatrix} s_j & t_j \\ s_{j+1} & t_{j+1} \end{bmatrix},$$

thus in step 5 if $\deg q_j > \deg q_{\max}$ then $s_{\max}$ and $t_{\max}$ are easily update by the entries of the first row of R_j. Again by induction, in step 6 $q^*_{\max}$ represents the quotient with maximal degree in $\{q_{j+1}, \ldots, q_h\}$. But $s^*_{\max}$ and $t^*_{\max}$ are not on the same row as $q^*_{\max}$ in the Euclidean algorithm for r_0 and r_1. Let l represent the index of $q^*_{\max}$ in the EEA for r_0 and r_1. In step 6, if $\deg q^*_{\max} > \deg q_{\max}$ then we require to update $s_{\max}$ and $t_{\max}$ by s_l and t_l, respectively. According to the definition of R_l we have

$$\begin{aligned} \begin{bmatrix} s_l & t_l \\ s_{l+1} & t_{l+1} \end{bmatrix} &= R_l = Q_l Q_{l-1} \ldots Q_{j+1} R_j \\ &= \begin{bmatrix} s^*_{\max} & t^*_{\max} \\ m_1 & m_2 \end{bmatrix} R_j, \end{aligned}$$

where $m_1, m_2 \in F[x]$, hence

$$\begin{bmatrix} s_l & t_l \end{bmatrix} = \begin{bmatrix} s^*_{\max} & t^*_{\max} \end{bmatrix} R_j.$$

So to update $s_{\max}$ and $t_{\max}$ we simply multiply the vector $\begin{bmatrix} s^*_{\max}, t^*_{\max} \end{bmatrix}$ by matrix R_j. Therefore, at the end of step 6, $q_{\max}$ holds the quotient with maximal degree in $\{q_1, \ldots, q_h\}$ and $s_{\max}$ and $t_{\max}$ have the same index as $q_{\max}$ in the EEA for r_0 and r_1. This implies that the final results in step 8 are correct. Note that the EEA should be modified as well to return the maximal quotient and the corresponding values of s and t in step 1. We now show how to call MFEEA to compute the desired rational function.

Fast Maximal Quotient RFR Algorithm(FMQRFR)
Input: $f, g \in \mathbb{Z}_p[x]$ with $g \neq 0$, $\deg f > \deg g \geq 0$, and $T \in \mathbb{N}$
Output: Either $n, d \in \mathbb{Z}_p[x]$ satisfying $n/d \equiv g \bmod f$, $\mathrm{lc}(d) = 1$, $\gcd(n,d) = 1$, and $\deg n + \deg d + T < \deg f$, or FAIL implying no solution exists

1. $r_0 = f/\mathrm{lc}(f), \quad r_1 = g/\mathrm{lc}(g)$
2. $h, \rho_{h+1}, R_h, q, \tilde{s}, \tilde{t} = \mathrm{MFEEA}(r_0, r_1, \deg r_0)$
   if $\deg q \leq T$ then return FAIL
3. $\tilde{r} = \tilde{s} r_0 + \tilde{t} r_1$
   if $\gcd(\tilde{r}, \tilde{t}) \neq 1$ then return FAIL
4. $n = \mathrm{lc}(g)/\mathrm{lc}(\tilde{t}) \cdot \tilde{r}$
   $d = 1/\mathrm{lc}(\tilde{t}) \cdot \tilde{t}$
   return (n, d)

As pointed out earlier r is obtained from s and t using $r = sf + tg$, but $\tilde{s}$ and $\tilde{t}$ that are returned as the corresponding values of q, the quotient with maximal degree, are off by a constant factor. From the definitions of s and t we find that $s = \tilde{s}/\mathrm{lc}(f)$ and $t = \tilde{t}/\mathrm{lc}(g)$ and hence

$$\frac{r}{t} = \frac{\dfrac{\tilde{s}}{\mathrm{lc}(f)} f + \dfrac{\tilde{t}}{\mathrm{lc}(g)} g}{\dfrac{\tilde{t}}{\mathrm{lc}(g)}} = \frac{\mathrm{lc}(g)(\tilde{s} r_0 + \tilde{t} r_1)}{\tilde{t}} = \mathrm{lc}(g) \cdot \frac{\tilde{r}}{\tilde{t}}.$$

If we let $m = \deg f$, then step 2 takes $O(\mathsf{M}(m) \log m)$ operations in $\mathbb{Z}_p$. To compute $\tilde{r}$ in step 3, we perform two multiplications on polynomials of size at most m and one addition. The total cost for computing $\tilde{r}$ is thus $2\mathsf{M}(m) + O(2m)$ operations in $\mathbb{Z}_p$. Checking the coprimality of $\tilde{r}$ and $\tilde{t}$, using the FEEA, takes $O(\mathsf{M}(m) \log m)$ operations in $\mathbb{Z}_p$. Steps 1 and 4 both cost $O(m)$ operations in $\mathbb{Z}_p$. Thus the asymptotic cost of the algorithm is $O(\mathsf{M}(m) \log m)$.

The following algorithm is an extension of Wang's algorithm for $F[x]$ and uses the FEEA instead of the EEA.

Wang's Fast Rational Function Reconstruction Algorithm
Input: $f, g \in F[x]$ with F a field, $g \neq 0$ and $M = \deg f > \deg g \geq 0$
Output: Either $n, d \in F[x]$ satisfying $n/d \equiv g \bmod f$, $\mathrm{lc}(d) = 1$, $\gcd(n,d) = 1$ and $\deg n + \deg d < M$, or FAIL implying no such n/d exists

1. $N = \lfloor M/2 \rfloor, \ D = M - N - 1$
   $r_0 = f/\mathrm{lc}(f)$, $t_0 = 0$
   $r_1 = g/\mathrm{lc}(g)$, $t_1 = 1$
2. $h, \rho_{h+1}, R_h = \mathrm{FEEA}(r_0, r_1, \deg r_0 - N - 1)$
3. $n = r_{h+1} = s_{h+1} r_0 + t_{h+1} r_1$
   $d = t_{h+1}$
   if $\gcd(n,d) \neq 1$ then return FAIL
4. $n = \mathrm{lc}(g)/\mathrm{lc}(d) \cdot n$
   $d = 1/\mathrm{lc}(d) \cdot d$
   return (n, d)

If the FEEA is also used for computing $\gcd(n,d)$ in step 3, then the time complexity of Wang's algorithm would be $O(\mathsf{M}(M) \log M)$ as well. Algorithm FMQRFR normally must compute all the quotients to determine the largest but Wang's algorithm stops half way, and hence, is expected to take half the time. On the other hand Wang's algorithm outputs n/d only if

$$\deg f \geq 2 \max(\deg n + \deg d).$$

But the Maximal Quotient algorithm requires

$$\deg f > \deg n + \deg d + T,$$

which requires only one more point than the minimum necessary when T is chosen to be 1. The following table compares the time of both algorithms where the timings in columns 2 and 3 are for the EEA and the timings in columns 4 and 5 are for the FEEA. We have chosen n/d and f with $\deg n = \deg d$ and $\deg n + \deg d + 2 = \deg f$ where the coefficients of f, n, d are chosen at random from $\mathbb{Z}_p$. The data shows that Wang's algorithm (both versions) is almost 2 times faster than the maximal quotient algorithm (both versions). All the timings are in milliseconds.

deg f	MQRFR	Wang	FMQRFR	Fast Wang
64	21.72	12.50	21.51	12.37
128	80.94	38.45	83.54	169.01
256	278.18	142.93	212.04	131.65
512	1012.35	531.35	681.13	332.43
1024	3874.80	1903.78	1954.20	1056.44
2048	14802.39	7185.59	6325.23	3319.55
4096	59086.61	28345.89	20664.24	10875.93

5. OPEN PROBLEMS

Let $f, g \in \mathbb{Z}_p[x]$ where $f = \prod_{i=1}^{n}(x - \alpha_i)$, $n = \deg f > \deg g \geq 0$ and p is prime. Let q be a quotient in the EEA for inputs f, g and $k \in \mathbb{N}\backslash\{1\}$. If $\alpha_i \in \mathbb{Z}_p$ is chosen uniformly at random and g is a random polynomial, then

$$\mathrm{Prob}(\deg q \geq k) \simeq \frac{n-k+1}{p^{k-1}}.$$

The above conjecture implies that in the EEA the probability of getting a quotient of degree 2 or more is almost equal to n/p. Therefore if p is large enough compared to n, then with high probability all quotients are of degree less than 2. Hence in the Fast Maximal Quotient RFR Algorithm by choosing $T = 1$, with high probability we get a correct result and we require f to be at least of degree $\deg n + \deg d + 1$.

6. REFERENCES

[1] Richard P. Brent, Fred G. Gustavson, and David Y. Y. Yun. Fast solution of Toeplitz systems of equations and computation of Padé approximants. *Journal of Algorithms*, 1:259–295, 1980.

[2] G. E. Collins and M. J. Encarnacion. Efficient Rational Number Reconstruction. *J. Symbolic Computation* **20**, pp. 287–297, 1995.

[3] M. J. Encarnacion, Computing GCDs of Polynomials over Algebraic Number Fields, *J. Symbolic Computation* **20** (1995), pp. 299–313.

[4] Sara Khodadad. Fast Rational Function Reconstruction. Master's thesis, Simon Fraser University (SFU), Burnaby, BC, Canada, 2005.

[5] L. Langemyr, S. McCallum, The Computation of Polynomial GCD's over an Algebraic Number Field, *J. Symbolic Computation* **8** (1989), pp. 429–448.

[6] Daniel Lichtblau. Half-gcd and fast rational recovery. In *ISSAC '05: Proceedings of the 2005 International Symposium on Symbolic and Algebraic Computation*, pages 231–236. ACM Press: New York, NY, 2005.

[7] R. E. Maeder. Storage allocation for the Karatsubs integer multiplication algorithm. In *DISCO '93: Proceedings of the Interational Symposium on the Design and Implementation of Symbolic Computation Systems*, Springer-Verlag LNCS, 1:59–65, 1993.

[8] R. T. Moenck. Fast computation of gcds. In *STOC '73: Proceedings of the fifth annual ACM Symposium on Theory of Computing*, pages 142–151. ACM Press: New York, NY, 1973.

[9] Michael Monagan. Maximal quotient rational reconstruction: An almost optimal algorithm for rational reconstruction. In *Proceedings of the International Symposium on Symbolic and Algebraic Computation*, pages 243–249. ACM Press: New York, NY, 2004.

[10] Peter Lawrence Montgomery. *An FFT extension of the elliptic curve method of factorization*. PhD thesis, Los Angeles, CA, USA, 1992.

[11] Victor Y. Pan and Xinmao Wang. Acceleration of Euclidean Algorithm and Extensions. In *Proceedings of the International Conference on Symbolic and Algebraic Computation*, pages 207–213. ACM Press: New York, NY, 2002.

[12] A. Schönhage. Schnelle Berechnung von Kettenbruchentwicklungen. *Acta Informatica*, 1:139–144, 1971.

[13] Allan Steel. Private Communication.

[14] Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. Cambridge University Press: Cambridge, New York, Port Melbourne, Madrid, Cape Town, second edition, 2003.

[15] Paul S. Wang. A p-adic Algorithm for Univariate Partial Fractions. In *Proceedings of the fourth ACM Symposium on Symbolic and Algebraic Computation*, pages 212–217. ACM Press: New York, NY, 1981.