

# Rational Function Computing with Poles and Residues

Richard J. Fateman  
Computer Science Division, EECS  
University of California, Berkeley

January 10, 2013

## Abstract

Computer algebra systems (CAS) usually support computation with exact or approximate rational functions stored as ratios of polynomials in “expanded form” with explicit coefficients. We examine the consequences of introducing a partial-fraction type of form in which some of the usual rational operations can be implemented in substantially faster times. In this form an expression in one variable say  $x$  is expressed as a polynomial in  $x$  plus a sum of terms each of which has a denominator  $x - c$  perhaps to an integer power, where  $c$  is in general a complex constant. We show that some common operations including rational function addition, multiplication, and matrix determinant calculation can be performed many times faster than in the conventional representation. Polynomial GCD operations, the costliest part of rational additions, are entirely eliminated. Application of Cauchy’s integral theorem allow for trivial integration of an expression around a closed contour. In some cases the approximate evaluation of transcendental functions can be accelerated, especially in parallel, by evaluation of a formula in pole+residue form.

## 1 Introduction

Computer algebra systems (CAS) usually support computation with exact rational functions stored as ratios of polynomials in expanded form with explicit coefficients. For example,

$$\frac{7x^3 - 70x^2 + 231x - 252}{x^2 - 11x + 30}. \quad (1)$$

Alternative forms are sometimes used for particular algorithms or for attaining special efficiencies. One of these is factored (or partly factored) form:

$$\frac{7(x-4)(x-3)^2}{(x-6)(x-5)}. \quad (2)$$

Another representation rearranges polynomials using Horner’s rule:

$$\frac{x(x(7x-70)+231)-252}{(x-11)x+30}. \quad (3)$$

It is also useful to see approximations (which in some cases turn out to be exact) such as a Taylor series expansion, this one expanded at  $x = 0$ :

$$-\frac{42}{5} + \frac{231x}{50} - \frac{539x^2}{1500} - \frac{2359x^3}{45000} + \dots \quad (4)$$

A less familiar variation of representation is the use of “straight-line programs” (SLP). One possible SLP version might be:

```
f(x) :=
{t1=x-3; t2=t1*t1; t3=x-4; t4=t2*t3; t5=7*t5;
 t6=x-5; t7=x-6; t8=t7*t6; return(t5/t8)}
```

A “black-box” version might be just like this SLP or some equivalent algorithm, but with the detailed computation hidden. The only information available from a black-box version is a response, for each input value, of a corresponding output value. It is surprising how many computations can be performed (fast) on this apparently impoverished representation [6].

The form that is provided by most CAS that is closest to the topic of this paper is the partial-fraction expansion form:

$$-\frac{28}{x-5} + \frac{126}{x-6} + 7x + 7. \quad (5)$$

While CAS can convert to this form, they usually do not support efficient further computation maintaining this form, especially through multiplication operations. This need not be the case, since some operations are easily programmed in this form. Adding two partial fraction forms is simple. e.g. to add  $1/(x-7)$  to the expression above requires only a merge to get

$$-\frac{28}{x-5} + \frac{126}{x-6} + \frac{1}{x-7} + 7x + 7. \quad (6)$$

To multiply by  $1/(x-7)$  yields

$$\frac{14}{x-5} + \frac{126}{x-6} + \frac{168}{x-7} + 7. \quad (7)$$

Versions of this last representation, giving explicit poles and residues are frequently encountered in digital filter design. We have not yet found a situation justifying *large-scale* direct applications here for semi-symbolic methods, but such applications may be possible. In particular, it seems to us that the typical representation of functions via polynomial coefficients is usually indirect from physical measurements. It is more plausible that measurements would be interpolation points which could be used to derive a polynomial. For rational functions, physical measurements might reasonably derived from locations of poles<sup>1</sup>. If we step beyond the question of where data for these problems originate, what can we say?

Conversion of rational expressions into this format is commonplace for integration and Laplace transforms. Most CAS must therefore provide it,

The question for us in this paper is whether maintaining this format internally through sequences of operations is useful.

We cannot say if it is useful without some context, but we can address the question of how to design algorithms that are efficient in time and space, and whether it is convenient in other respects.

Historically, the lack of support means that there have been few investigations. It can be helpful in elucidating the potential for numerical overflow when evaluating at a point near the zeros of a denominator. It can also be useful for computations consisting of sequences of multiplies and adds; for example a matrix determinant. (Note that there are popular methods for computing determinants that require division. Such methods are a bad idea in this representation. Fortunately some of the best *symbolic* methods do not use division [10].)

Barton Willis, independent of this work, requested assistance<sup>2</sup> in algorithms for fast manipulation of partial-fraction forms—namely preserving the partial fractions, and therefore *not* dependent on polynomials zero-finding. His application (unpublished) involves the solution of differential equations in terms of  ${}_2F_1$  hypergeometric functions, generalizing work by Bronstein and Lafaille [2].

We note that “renaming” the denominators as new indeterminates can convert such rational functions into multivariate polynomials, and this provides some of the same speed-up in some operations. Renaming must be accompanied by the simplification aspects (cancellation) of these new indeterminates. This consideration is essentially what our representation expresses.

<sup>1</sup>Equations 5,6,7 are somewhat deceptive. In more realistic examples those simple integers would be more likely to be floating-point complex numbers representing approximations of locations of poles in the complex plane and residues.

<sup>2</sup>Maxima mailing list, Jan. 19, 2005

## 2 Poles, Residues, Zeros

For the rest of this paper we discuss adopting a version of partial fractions, but with the additional requirement that all factors of the denominator be monic linear factors (again, as in the example). The usual partial-fraction expansion program in a computer algebra system may compute “over the rationals” and thus may not allow a complete decomposition of terms with irreducible (over the rationals) denominators of  $1/(x^2 + 1)$ . Our requirement of linear factors means that we need a partial-fraction expansion over algebraic numbers, or more likely, approximations of algebraic numbers. Thus we would express  $1/(x^2 + 1)$  as  $0.5i/(x + i) - 0.5i/(x - i)$ , or in reality as a compact data structure storing pairs of complex floating-point numbers,  $(0.5i, 1.0i)$ ,  $(-0.5i, -1.0i)$ : each pair has a residue and a root.

In general, producing an expansion requires finding approximations to roots of polynomials in the complex plane, an operation available in most scientific subroutine libraries and CAS. Maple in fact provides conversion of an expression  $E$  to such a form in *one* built-in command, `convert(E,parfrac,x,complex)`. This conversion is not usually reversible since rational approximations for roots cannot always be algebraically faithful. `normal(convert(1/(x^2+3)^2,parfrac,x,complex))`

Before proceeding further, we wish to anticipate some objections to our proposal:

- The representation we propose does not generalize to more than one variable *appearing in the denominator*. It allows extra parameters in numerators. The presence of multiple variables in denominators leads to computational difficulty. For example we would need a good way of representing the “pole” in  $f(x, y) := 1/(x - y)$ . Considering how many papers are written on algorithms that are only good for *polynomials* in one variable, a much more restricted domain, perhaps we should not worry about this! Rational functions in one variable provides more coverage.
- There are possible problems with numerical stability: the traditional objective in computer algebra systems (CAS) is to get exact answers when possible. Here, except for artificially-constructed examples with easily found exact rational roots, we expect to commit, in the first step, an approximation for algebraic numbers by (in general) complex floating-point numbers. We need these in order to represent polynomial roots which are likely to be at non-rational locations. Of course the initial statement of the problem may be approximate already— indeed, it may be that the location of poles or zeros is a measured property rather than the coefficients of a polynomial.

Should the problem originate as a polynomial, A typical CAS allows for root-finding using hardware floating-point and generally software arbitrary precision as well. It is difficult to anticipate how precise a computation might be needed at the outset, and sensitivity of the root-finding to distinguish multiple zeros (that is, higher-order poles), from a cluster of nearby zeros is a concern.

A path that can be pursued in the context of computer algebra systems is to *remove multiple zeros by exact square-free factorization*, an inexpensive operation. This speeds up the root-finding and cleans up the results, but still leaves open the possibility that nearly equal roots originating from different inputs will have to be judged as being distinct or not. In particular the complexity of the arithmetic may depend on matching the location of poles from such different inputs. If the algorithm decides that two barely-distinct poles are instances of the same pole, such a collapse of two into a multiple pole reduces the size of the representation and simplifies computation<sup>3</sup>.

What is worth noting though, is that the poles are present in the input to the addition, multiplication, and determinant programs, *and they determine exactly the poles that may be present in the answer*. That is, no poles are created, nor are additional errors in their locations introduced by roundoff. It is possible that the order of the poles will be increased (e.g. by squaring) or decreased, perhaps to order zero. If the arithmetic is done using floating-point, there is a possible accumulation of errors in the numerators. While these are potential problems, there also is a positive side: In some examples it is possible to use the approximate answers to derive exact results: by bracketing known integer-valued problems so that only one integer is in an approximating interval. The approximate calculations necessary to do this can be far faster than attempting the exact calculation directly.

---

<sup>3</sup>Surprisingly, this does not so much affect some aspects of computing, such as the accuracy of evaluation of the expression away from the pole location.

In passing, we note that the traditionalists among CAS designers expect them to deal only with integer-valued or exact rational coefficients in the representation: yet what physical measurement concludes by returning a polynomial of degree 10,000 with exact integer coefficients? Is this more or less plausible than returning a rational function represented by the residues at its (perhaps approximate) zeros?

- Conversion from conventional to pole/ residue form tends to produce numbers that look “messy” to humans. For computer internal representations it takes the same space and uses the same computation time in dealing with a double-precision number regardless of its aesthetic appeal. That is, a number that can be printed accurately using only a few digits, e.g. 12.3 is no more efficient than one that can require 16 digits and an exponent. This is typically what we see in computing approximations to roots of polynomials. That is, an apparent disadvantage *to human consumption* of the pole/residue representation is that it is more likely that the forms may seem to have “lots of digits”.

As an example of numerical difficulties that might occur during the translation of a conventional form into a residue/pole form, consider the simplest case of a double pole at zero in  $f(x) = 1/x^2$ . Programs will likely find these roots exactly but let us play devil’s advocate and (mis)represent the roots instead by a pair of equally spaced poles  $\pm\epsilon$  away from zero. Then we have a function

$$g(x) := \frac{1}{x^2 - \epsilon^2} = \frac{1/(2\epsilon)}{x - \epsilon} - \frac{1/(2\epsilon)}{x + \epsilon}.$$

While the two forms are mathematically identical, the second (partial-fraction) version of  $g(x)$  is computationally less accurate where  $|x| \gg 0$  since (as written) it would be computed as the difference of two nearly-equal values.

Mathematically,  $g(x)$  is very close to  $f(x)$  except near  $x = 0$ . In that case we can compare the definition of  $f$  to the Taylor series for  $g$ :

$$f(x) := 1/x^2$$

$$g(x) := -\frac{1}{\epsilon^2} - \frac{x^2}{\epsilon^4} - \frac{x^4}{\epsilon^6} + \dots$$

How close are these? The worst relative error is when  $x = 0$ . At that point the approximation gives  $1/\epsilon^2$  instead of  $\infty$ . Nearby, say with  $\epsilon = 10^{-8}$ , the result,  $g(10^{-10}) = 1.0001 \times 10^{20}$  can be compared to the exact answer,  $f(10^{-10}) = 1.0000 \times 10^{20}$ .

The point here is that the functions, while not mathematically identical, are computationally difficult to distinguish by plotting, and may serve identically in ordinary engineering design. To those who object to any error at all or who insist that the computation in question has all exact numerical parameters and that one must always correctly distinguish between a cluster of  $N$  nearby zeros versus a zero of order  $N$ , this may be simply an inappropriate representation.

We have not attempted to completely characterize the efficiency or the nature of the errors produced by polynomial root-finding programs [12, 14]. While we suggest how one might *force* clusters of zeros into one multiple zero by reducing precision and rounding (see Appendix), we are not promoting it within the operations. Any of these errors in the denominators are, arguably, committed *prior* to using our programs. Error propagation (roundoff) committed in evaluation at those points during multiplication or addition of numerators will happen as ordinary floating-point arithmetic: various rearrangements of computations can reduce this effect, but not eliminate it. We reiterate: the poles remain in the same locations.

Nevertheless, we note that the cost of such conversion from more traditional representations as well as the nature of errors in a modern CAS setting may be highly variable: it may be plausible to base an algorithm on finding approximations to extremely high accuracy by means of multiple-precision (software) floating-point arithmetic. If the original coefficients in the problem have some possible error then these input errors can be made to dominate the calculation<sup>4</sup>.

It is well known that given all poles, (with multiplicities), and their residues plus a “direct term” (polynomial part) we can determine a unique rational function. Perhaps the most direct algorithmic assertion

---

<sup>4</sup>Some of the historical concern for avoiding approximations in CAS seems to have faded, judging by the number of recent publications on approximate polynomial GCD, including some more “applied” researchers (consider [15, 8, 5]). In reality we have not seen any version of an approximate polynomial GCD algorithm become a central part of a computer algebra system, although Maple now has a SNAP [17] library with some approximate GCD algorithms.

of this is in the Matlab system [13] which provides a program for changing between two rational function representations; we take this as a partial confirmation that this kind of representation is of interest in some applications. Oddly enough, Matlab syntax here uses the same name for a function and its inverse:

```
[r,p,k] = residue(b,a)
[b,a]   = residue(r,p,k)
```

The operation `residue(b,a)` takes a pair of polynomials  $\{b, a\}$  and returns a triple, say  $[r, p, k]$ , which is equivalent to  $b/a$  in “pole-residue” representation. This triple is a list of residues, poles, and a (polynomial) direct term, computed as part of a partial-fraction expansion for  $b/a$ . The polynomials  $b$ ,  $a$  and  $k$  (in some unstated parameter, say  $s$ ) are represented as lists. For example,  $[3,4,5]$  is  $3s^2 + 4s + 5$ . These Matlab conversions are done with complex floating-point numbers and are naturally subject to numerical error if the polynomial roots are not representable exactly, or if errors are committed in the solving process.

*Matlab does not support operations in the pole-residue form, however. We explore this possibility in the remainder of this paper.*

### 3 Arithmetic with poles and residues

## 4 Conversion from ratio of polynomials to poles/residues form

Let us convert a ratio of polynomials  $M$  in the variable  $x$  into a poles/residues form

Step 1:

Divide, with remainder, the numerator of  $M$  by the denominator  $D$ .

The quotient  $Q$  is a polynomial which becomes the polynomial part of the result. Call the remainder  $N$ . We must now find the full partial fraction representation of  $N/D$ .

Step 2: Factor over the complex numbers the denominator  $D$ .

This can only be done in general approximately by numerical root finding.

Step 3: Let  $DR$  be the number of distinct roots,  $c[1], \dots, c[DR]$ . Let  $MUL[k]$  be the multiplicity of root  $k$ ,  $1 \leq k \leq DR$ .

Step 4: Let  $qq[k] =$  the product of  $(x-c[j])^{MUL[j]}$ ,  $1 \leq j \leq DR$   $j \neq k$ .

Define  $r[c[k],j]$ :

```
let f= N/qq[k]
```

```
also let h=MUL[k]-j.
```

```
If h=0 then return the result of substituting c[k] for x in f.
```

```
otherwise
```

```
return the result of substituting c[k] for x in 1/j!*diff(f,x,h),
```

```
where diff(f,x,h) is the hth derivative with respect to x of f.
```

The partial fraction expansion will then consist of  $Q$  plus the sum of terms for each  $k= 1 \dots DR$

and for each multiplicity  $j=1..MUL[k]$  of  $r[c[k],j]/(x-c[k])^j$ .

An implementation of this in the CAS Maxima is in the appendix.

## 4.1 Addition

Adding two rational functions in residue form is easily done in linear time:  $A = r_1/(s - p_1) + r_2/(s - p_2)$  can be added to  $B = r_3/(s - p_3) + r_4/(s - p_4)$  if the poles  $\{p_i\}$  are distinct, just by merging lists of poles and residues. The “definite” parts are added as polynomials. The data structure holding the poles and the collection of residues at that position must allow for rapid search and insertion. During the merge, if two terms have the same pole, the residues must be added, and if their sum is zero, the term must be removed. Our programs actually use hash tables.

To have any hope of a canonical external appearance, the lists of poles should be uniquely ordered. The poles may be sorted by a lexical ordering, perhaps real then imaginary parts. The sorting need not be done except on output, and so this is not actually part of the time for addition.

The addition operation can be performed in parallel *in constant time* if you have enough processors. Consider all the work devoted to speeding up the polynomial GCD algorithm that is totally avoided here.

## 4.2 Evaluation

We deal with evaluation before multiplication because this operation is required in the multiplication of expressions, at least as we propose in the next section, where multiplication at a pole point  $p$  requires the evaluation of the (other) expression at that point  $p$ . While there are a variety of somewhat faster techniques for evaluating a polynomial at a point, or a collection of points, involving preconditioning or other rearrangements, Horner’s rule is typically acceptably fast and requires no preconditioning. Consider the cost then of an evaluation of a pole/residue expression at a generic (i.e. non-pole) point. The polynomial part, say of degree  $d$  can be done by Horner’s rule, in  $O(d)$  operations, assuming (falsely in general, but more nearly true for floating-point) that coefficient arithmetic is constant cost. For each of  $q$  isolated (non-multiple) poles, the cost is one subtraction to compute the denominator, and one division, for cost  $O(q)$ . If there is a multiple pole of order  $m$  then the cost is  $O(m)$  if we used Horner’s rule in the obvious way, computing a polynomial at a numerical point  $1/(x - p)$ . Thus the evaluation cost is essentially linear in the number of terms, either polynomial or residue/pole, given a plausible dense case expression. Another issue is numerical accuracy. If the point location is very nearly at a root, the denominator of that term will be nearly zero and so any non-zero numerator value will be magnified. If the root is inaccurately computed, this magnification will happen at another location (perhaps close to the root). This does not seem inherently more hazardous than evaluating a high-degree polynomial divisor more-or-less accurately by floating-point methods; it is more hazardous if the denominator can be evaluated exactly. That possibility exists if arbitrary-precision arithmetic is used, but represents a substantially more costly operation.

## 4.3 Evaluation of trigonometric (etc) functions.

Consider evaluation of  $\sin(x)$  by Taylor series to degree 13. An alternative to evaluating the series is to evaluate a rational function (Pade approximation) which agrees with the Taylor series (to degree 13), and looks like

$$\frac{479249x^7 - 52785432x^5 + 1640635920x^3 - 11511339840x}{18361x^6 + 3177720x^4 + 277920720x^2 + 11511339840}$$

which can be converted to a poles/residues expression we have formatted for humans viewing and typesetting, interspersing syntactic sugar showing denominators, as well as much abbreviated floating-point printing of coefficients:

```
3.0048d3/(x+9.0402*i)+
(1.4027d3*i+3.4565d2)/(x+8.1627*i+4.5779)+
(3.4565d2-1.4027d3*i)/(x+8.1627*i-4.5779)+
(3.4565d2-1.4027d3*i)/(x-8.1627*i+4.5779)+
(1.4027d3*i+3.4565d2)/(x-8.1627*i-4.5779)+
3.0048d3/(x-9.0402*i)-
(479249*x)/18361
```

An advantage to such an arrangement is that there is obvious parallelism. (We are not recommending this particular computation scheme for  $\sin(x)$ ; the cost of complex arithmetic and finally adding together the pieces may make it more expensive than other methods, regardless of parallelism. We chose this illustration to keep the discussion brief.

## 4.4 Multiplication

Multiplying two  $[\mathbf{r}, \mathbf{p}, \mathbf{k}]$  forms is somewhat messy. We are unaware of any prior explicit description of these algorithms in the literature. One way is to do essentially the computation for conversion from conventional ratio-of-polynomials form, but starting at STEP 3, bypassing the expensive part of root-finding. (See appendix).

Each of the inputs  $R, S$  has pole-terms and a polynomial (or “definite” term). By pole-term at  $p$  we mean the sum of all the residues of different order,  $\sum a_i/(x-p)^i$ .

If we use a classical approach to multiplications, we form the cross product of each term in  $R$  with each term in  $S$ , and add them all together.

Multiplication of the purely polynomial terms can be done by any standard algorithm, with corresponding complexity. The asymptotically fastest known methods are based on FFT techniques; it may be more reasonable to estimate the timing of small examples by assuming the classical methods, in which case the cost for multiplying  $n$  terms by  $m > n$  terms, including sorting terms into canonical order is  $O(d^2nm \log n)$  where  $d$  is related to a bound on the size of the coefficients. Multiplication of *matching* poles in  $R$  and in  $S$  is done by matching them by order and multiplying the residues. This is mapped fairly directly into polynomial multiplication (essentially of Laurent series), with corresponding complexity.

Two other cases remain: Multiplication of a polynomial by a pole-term, and multiplication of pole-terms at different poles.

Consider the first of these, pole-term times polynomial. The simplest, and we expect by far the most common case, is that of a simple pole, in which  $p(x) \times (a/(x-b))$  is computed by a polynomial division of  $p$  by  $x-b$ , returning quotient  $Q$  and remainder  $R$ . The result is  $aQ + aR/(x-b)$ . For multiple poles

$$\sum_{i=1}^m \frac{a_i}{(x-p)^i} \times \sum_{j=0}^n b_j x^j$$

Expand the second term (the polynomial) in a power series in  $(x-p)$ . The coefficients  $\{c_i\}$  can be computed by “synthetic division”. Also, let us re-index the  $\{a_i\}$  to correspond to negative numbers:

$$\sum_{i=-m}^1 a_i (x-p)^i \times \sum_{i=0}^n c_i (x-p)^i.$$

If we “virtually” multiply by  $(x-p)^m$  then the computation above is equivalent to a polynomial multiplication. The  $m$  lower order terms become the pole result. Thus the complexity is still dominated by the cost of a related polynomial multiplication.

The second case is multiplication of pole-terms that are not at the same pole. Our programs assume that any two poles with different floating-point representations are distinct. Any interest one might have in claiming that two (or more) poles that are “very close” should be perturbed to be at the same location are outside the scope of our arithmetic operations.

Some algebra (assisted by a CAS) led us to the following (so far as we know, new) formulas for generating products. The first formula computes an expansion for  $w_n := 1/((x-q)^n(x-p))$ :

$$w_n := \frac{1}{(p-q)^n(x-p)} - \sum_{i=1}^n \frac{1}{(p-q)^i(x-q)^{n+1-i}}$$

Note that the right-hand side explicitly displays the denominators of powers of  $(x-p)$  and  $(x-q)$ .

We can incorporate the formula into a more general result, for an expansion for  $k_{m,n} := 1/((x-q)^n(x-p)^m)$ :

$k_{m,n} := \text{if } m = 1 \text{ then } w_n \text{ else}$

$$\frac{1}{(p-q)^n (x-p)^m} - \sum_{i=1}^n \frac{k_{m-1,n-i+1}}{(p-q)^i}.$$

Note again that the right-hand side explicitly displays the denominators of powers of  $(x-p)$  and  $(x-q)$ .

In practice, a computation would consist of arranging (negative) powers of  $(p-q)$  and adding these residues into the correct slots indexed by  $p, q, n, m$ . Powers of  $(x-p)$  or  $(x-q)$  do not require arithmetical “computation” since the corresponding coefficients are merely indexed into slots. The summations in the formulas are not suggesting that we are adding anything, but in a similar fashion are denoting an indexed vector of the different pole-term components.

These formulas translate directly into a program for manipulating the data structures. There are  $O(mn)$  coefficient multiplies to produce  $n+m$  terms.

After programming the above formulas we found their performance to be disappointing, and so we re-thought the operation and programmed what appeared to be a much faster (perhaps new) version of multiplying pole-terms by pole-terms by essentially computing a truncated Laurent expansion (up to power -1) of such products around  $a$ , that is, the coefficients of  $1/(x-a)^n$  by evaluating the other factor, as well as sufficiently many of its derivatives, at the constant  $a$ . The expansion has to be computed about each of the two pole locations in the product. In particular, the product

$$f(x) \times g(x) = \left( \sum_{i=1}^{k_1} \frac{r_i}{(x-a)^i} \right) \times \left( \sum_{i=1}^{k_2} \frac{s_i}{(x-b)^i} \right)$$

can be computed by considering the two sets of pole terms, expansion about  $a$  and expansion about  $b$ . Let

$$g^{(n)}(a) := \frac{1}{n!} \left. \frac{d^n g(x)}{dx^n} \right|_{x=a},$$

and  $g^{(0)}(a) = g(a)$ . Then the terms at  $a$  are

$$\sum_{i=0}^{k_1-1} \frac{\sum_{0 \leq j \leq i} r_{n+j-1} g^{(j)}(a)}{(x-a)^{n-i}},$$

and symmetrically around  $b$ . These terms account for all the residues at the two locations. There are various “asymptotically fast” arguments that could be made for some of the sub-algorithms including evaluation of a function and its derivatives. However, using the obvious methods and characterizing each of the inputs as a worst case of  $t$  distinct poles gives us an  $O(t^2)$  multiplication (assuming constant coefficient arithmetic costs). This would be comparable to the multiplication of two ordinary rational functions whose numerator and denominator were of length  $t$ , except that the advantage of this program, compared to the usual CAS approach, is that polynomial GCDs are entirely avoided. The precise time advantage for this program therefore may depend primarily on the savings in avoiding whatever polynomial GCD method is used by the comparison methods. Historically, polynomial GCDs were generally considered a bottleneck<sup>5</sup>.

Further details on reducing the work in evaluating this form are given in the program listings, available from the author. We would encourage additional experimental evaluation in system contexts and tuning for other environments. A valid comparison using another system naturally requires implementing the pole-residue version in the other system.

By tradition one usually asks if an algorithm can be done in parallel: in this case, if one has multiple processors, one processor could be assigned to each distinct pole of one input to the multiplier, and one or more processors could be assigned to the polynomial part. There would be no synchronization necessary for the pole  $\times$  pole products, since the residue at a pole cannot depend on residues at other poles. The polynomial

<sup>5</sup>Exact univariate polynomial GCDs are no longer a terrible burden given the possibility of modular or sparse methods; the polynomials here are likely to have floating-point coefficients, complicating the matter. Nevertheless avoiding their computation entirely would seem to be a major benefit.



$\times$  pole product does require coordination, and the feasibility of parallel multiplication of polynomials has been well-explored in the literature.

Since some pole-product processes may require far more computation than others, it may be plausible to look for finer-grained parallelism. See also consideration of special hardware for the evaluation of a rational function at a point [9].

Note that the size of the answer may be exponential in the size of the input: e.g.  $(x^n + 1)/(x + 1)$  is of size  $O(\log n)$  but the expansion is  $O(n)$  terms long.

## 4.5 Division

We have no particularly attractive method for division  $R/S$  except to point out that if we succeed in computing the inverse  $1/S$ , then division reduces the problem to multiplication, once again.

We suggest converting  $S$  to a traditional form, swapping the numerator and denominator, (they are relatively prime) and finding the zeros and multiplicities of the new denominator. From this we can convert to residue/pole representation. Rootfinding is discussed further, below.

Some algorithms including matrix inversion use mostly multiplication and addition, but with occasional division; this process may be acceptably efficient in such cases.

## 4.6 Differentiation

This is a simple linear-time algorithm in which the polynomial part is differentiated, and for each pole term  $a(x - p)^{-n}$  we shift the exponent by one to produce a term  $-an(x - p)^{-n-1}$ .

## 4.7 Integration

For the rational part of a (line) integral, this is a simple linear-time algorithm in which the polynomial part is integrated, and for each pole term  $a(x - p)^{-n}$  for  $n > 1$  we shift the exponent  $(a/(n-1))/(x - p)^{-n-1}$ . For the cases  $a/(x - p)$  we must produce another form not in the residue/pole representation, namely a collection of  $a \log(x - p)$ . We suggest that these terms *not* be collected into a single logarithm via  $a \log(x - p) + b \log(x - q) = \log((x - p)^a (x - q)^b)$ , since  $a$  and  $b$  are in general not small integers. If they are, for example, complex floats, we do expect to have a polynomial argument to the log. Normally we would expect to see somewhat more comfortable forms for integration (avoiding full factorization unless needed, maybe using arctangents). Nonetheless, this is an effective procedure for integration, even if the representation's version of "simplest" does not correspond to the conventions typically used for small textbook problems intended to be done by hand. Some results would likely need to be massaged for comfortable human consumption.

It is also possible to consider contour integration. The knowledge of the poles and residues produces, via the Cauchy integral theorem, the contour integral on a closed curve by identifying the poles within the curve.

## 5 Finding zeros

Knowing only approximations to the zeros of denominators means that when we combine rational functions originating in different calculations, we may have differing approximations for values which, if computed exactly, would be the same. If it is necessary to "cancel" such terms, we may need to refer to more elaborate mechanisms to stage a reasonable algorithm for nearby matching [15]. It is possible, if one begins with exact coefficients, to find polynomial zeros to any desired accuracy, but for high accuracy this may require using software "bigfloat" packages. An appropriate choice of error bounds will depend on the application, but one plan might be to run the same program with increasingly higher accuracy until some level of confidence is attained. In some cases it may be possible to attempt reconstruction of exact rationals from (very close) floating-point numbers to confirm the exactness of calculations, or to use intervals. It is easy to see problems where irrational numbers are required and so any exact finite rational representation will necessarily be the value of a different more-or-less nearby quantity.

The cost for rootfinding for a polynomial of degree  $n$  to precision  $b$  bits has a lower bound of about  $O((n \log^2 n)(\log n + b))$  [16], though actual programs on realistical values of  $n$  may not exhibit such a low growth rate. McNamee’s bibliography on root finding [14], which has the important property of being updated periodically, includes far more references than can be reasonably cited here. We mention that there are several implementations for polynomial root approximation which deliver results to any requested accuracy. In conjunction with bounds on root separation, such programs can be used (given exact input coefficients) to determine definitively whether apparently “close” roots are actually identical or not. Rootfinding *per se* appears to be a problem solved many times over, even in the arbitrary-precision float domain. What is not resolved is how much additional accuracy might be advisable for obtaining an adequate result *after a sequence of divisions or conversions*. It should be clear that any algorithm requiring such a sequence might benefit from a transfer to a traditional representation as coefficient sequences. In passing, we mention that high-precision and high-accuracy root-finding programs also obviate the need to consider exponent overflow or underflow, since the representations used extend both precision and range. The recent interest in using floating-point approximations in (for example) polynomial GCD calculations may provide insight in particular applications.

## 6 Extensions, Limitations

Of particular interest (while contributing to significant additional pain!) are possible extensions to representations for functions of several variables (perhaps defined implicitly). We see no good way of doing this if more than one variable occurs in a general (non-homogeneous) setting in the denominator. Just to point out the problem, consider the rational function in two variables:

$$\frac{y^2 + 2y + x + 3}{y^3 + x^2(y^2 + 2y + 3) + 3y^2 + x(2y + 3) + 2y + x^3 + 1}.$$

The denominator is irreducible over the rationals, and we have no convenient way of finding “linear over the complex” roots of such expressions in general. On the other hand, in some specific cases where a rational function can be expressed (say) with a denominator that is homogeneous in some set of variables, we can consider computing a collection of linear factors as denominators. This kind of activity is shown in several papers at the SNAP [17] workshops, e.g. Bellido [1]. See also earlier work on the Durand-Kerner method [7].

For purposes of exposition we do not pursue this possibility but merely contrast it to the univariate situation in which we can always find such roots, (in principle, numerically). The above expression, evaluated at  $y = 0$  reduces to

$$\frac{1}{(x + 1)^2} + \frac{2}{(x + 1)^3}.$$

We know of no way to “lift” this solution to encompass the  $y$  dependency of the original rational function.

Instead of finding zeros, we could try factoring the denominators (in any number of variables, and perhaps over some algebraic extension field). This can provide a representation of rational functions as a sum in partial-fraction form, or as a ratio in which the denominator (and perhaps also the numerator) is factored. This is known as partially-factored form [3]. Since we cannot depend on multivariate denominators to be substantially factorable (and we don’t really view the construction of splitting fields as low cost or even necessarily advantageous) we might well see no positive effect. If the denominator doesn’t factor we end up with the usual representation: the ratio of two relatively prime polynomials.

We would also like rapid and accurate mappings between representations, especially including program forms in numerical languages.

Evaluating the usefulness of the interaction between the precision of root-finding programs and the appropriateness of this representation awaits some applications; this paper is intended to raise awareness of the representation and its consequences, not provide a complete catalog of applications.

## 7 Programs

Starting in November, 2003, we wrote a suite of programs for computing with univariate rational functions represented via poles and residues over “Common Lisp numbers” which include (generically), exact integer, exact rational, single- and double- precision floating-point numbers, and complex numbers (pairs of any of the previous number types). We can extend the operations to arbitrary-precision floats using one of several available packages in Lisp or by linking to libraries such as GMP/MPFR. When such floats are introduced, a user is immediately faced with a puzzle, namely how much precision should be specified.

We expect that once a calculation is initiated in which polynomial zero-finding programs are used, most applications will be compelled to continue subsequent calculations in complex-float format; perhaps initially double-precision if that is believed to be adequate. If overflow or loss of precision seems to present problems, an automated but pessimistic estimate of loss of precision could be computed via interval arithmetic, and the precision increased.

Also included in the suite of programs is a (Lisp) version of the Jenkins-Traub (ACM 419) zerofinding algorithm and a determinant expansion-by-minors program.

The operations mentioned above, as well as a few testing programs, including output routines are written in portable ANSI Common Lisp, and are available from the author.

### 7.1 Benchmarks

It is not our intention in this section to exhaustively compare programs on a wide, graded range of problems. It is our purpose to point out (essentially to an audience of designers of computer algebra systems or perhaps programmers intent on doing rational function arithmetic) that these methods have worked in at least one environment, ours. Others can obtain our code and/or reprogram our algorithms and see if favorable results can be obtained on their systems. We believe that some of the results are *so extremely advantageous in some computations* that an advantage on any system is highly likely to be visible. This requires that multiplication and (especially) addition are heavily used.

We obtained these figures on a 2.59GHz Pentium 4 computer, with programs written in Common Lisp (Allegro 6.2). Consider the time to multiply

$$\left( \sum_{1 \leq i \leq 40} \frac{i}{(x-10)^i} \right) \times \left( \sum_{1 \leq i \leq 40} \frac{i}{(x-20)^i} \right).$$

Viewing each of the inputs as a pole/residue form, we obtain the exact product after about 0.54 seconds of computation. This is performing all arithmetic exactly as rational numbers<sup>6</sup> We believe that this number can be reduced by carefully recoding or using a more efficient arbitrary-precision division. Profiling shows the division, in our environment, consumes about 85% of the time. If we instead feed in machine double-precision floats, the computation takes less than 0.047 seconds. If we insist *ab initio* that the program will work ONLY with double-precision floats, we can compile that into the program via declarations. A simple change of the arithmetic in one place – doing evaluation at a pole – from generic (which in Common Lisp allows for any combination of exact rationals, double-precision complex, and many alternatives in between), to *only double-precision floating point operations*, provided an additional speedup of about a factor of 8, or a time of about 0.006 seconds. Further speedups can be expected by changing other operations and converting the arrays to packed floats instead of “pointers to floats” as is the default in Lisp, or using one of the Lisp systems with a compiler that compiled to tighter machine arithmetic code.

Compare this to the canonical “ratio of polynomials” form in a similar computational environment. When done exactly in Maxima, the multiplication takes about 0.93 seconds, already slightly exceeding our current time (and far exceeding the 0.006 seconds for a specialized double-precision version). Providing the answer in partial fraction form takes an additional 7.5 seconds.

Adding the same two expressions in our poles/residues form provides a far more startling comparison. The addition takes a time that is far too small for our timer resolution; by executing it in a loop we find the

---

<sup>6</sup>The underlying implementation does not have very fast bignum or even floating-point arithmetic, so all times should be considered relative to this issue.

time to be about 0.00001 seconds. Adding the same two expressions in Maxima takes about 0.45 seconds, **about 45,000 times slower**.

As indicated earlier, we include in our suite of programs a version of RPOLY, ACM algorithm 419, (1972) by Jenkins and Traub. This polynomial zero finder, (code adapted from the Maxima computer algebra system) is an old standard, but still used widely. For finding double-precision roots of a well-behaved degree 40 polynomial, the program typically takes about 15 ms. On the other hand, if the zeros have some unfortunate configuration such as many multiple zeros, the time may be much higher to allow the program to refine the zeros; it may also fail to find zeros if the refinement fails. A number of recent alternative zero-finders can be considered.

We have experimented with assuming that the denominator of  $p/q$  is given exactly, and then removing all multiple zeros of  $q$  by dividing by the gcd of  $q$  and  $q'$ , its derivative<sup>7</sup>. This exact calculation then allows the zero-finder to assume that all zeros are distinct. Any intimation that there are multiple zeros means the precision should be increased. This can be done with any of several library programs for arbitrary-precision zero-finders. Further discussion of the benchmarking of zero-finders is beyond the scope of this paper.

We have set up some benchmarks for determinants, considering that some methods (not all) can be done with addition and multiplication only. We do so with some trepidation, because readers of this paper may trot out other ways of computing determinants and ask us to test them. Indeed, the Maple determinant program has a large number of option flags to set, which affect speed on certain examples.

We claim that, compared with traditional rational functions in fully canonical form, the pole/residue representation can be computed far faster, and that programs such as Maple might benefit from considering this representation.

First we discuss the appearance of the answer in this pole/residue form. The result naturally displays the collection of complex residues over poles of different orders. This may or may not be as compact as one that is a ratio of polynomials. In the typical cooked-up benchmarks for determinants, the best form for the answer is only obtained by factorization. Experimentation with “Cauchy matrix” forms is simple to set up and the results show a neat form. Take expressions from a class of matrices that look like this:

$$\begin{pmatrix} \frac{1}{2} & \frac{1}{x} & \frac{1}{x-1} \\ \frac{1}{x} & \frac{1}{2x-2} & \frac{1}{2x-3} \\ \frac{1}{x-1} & \frac{1}{2x-3} & \frac{1}{2x-4} \end{pmatrix}.$$

Determinants in this class can be computed at some expense as

$$\frac{x^3 - 8x^2 + 21x - 18}{32x^7 - 192x^6 + 456x^5 - 536x^4 + 312x^3 - 72x^2}$$

or at higher expense (requiring factoring) in this form:

$$\frac{(x-3)^2(x-2)}{8(x-1)^3x^2(2x-3)^2}$$

and at *very low expense* in the pole/residue form, which looks like this:

$$\frac{0.79167}{x} + \frac{0.25}{x^2} - \frac{2.125}{x-1} - \frac{0.5}{(x-1)^3} + \frac{1.33333}{x-1.5} - \frac{0.125}{(x-1.5)^2}$$

Converting this expression to conventional form, but leaving the denominator factored (and in general chopping some small coefficients) results in this smaller display:

$$\frac{0.125x^3 - x^2 + 2.625x - 2.25}{(x-1)^3x^2(2x-3)^2}$$

Turning from the “display format” issue to serious timing matters, we consider a different example. Define matrices with entries of the form  $g_{i,j} = \frac{1}{(x+j+i-1)^j}$ . We computed the determinant of the  $8 \times 8$  matrix using

<sup>7</sup>It is possible to construct a quick probabilistic modular test for “free of multiple roots.”

minor expansion [10] on a 2.59GHz Pentium 4 computer, using an implementation in Macsyma running on Allegro Common Lisp.

A typical run of the built-in program, (NOT using poles/residues) takes 262 seconds (18 in Lisp garbage collection).

Using EXACT poles/residues and the minor expansion method, this example uses 61+10=71 sec, of which 92 percent is bignum arithmetic. We have found that using GMP arithmetic is, for bignums larger than a few words, substantially faster than the built-in arithmetic in our benchmark Lisp, and thus we could expect a factor of 2-3 improvement by using the GMP library here.

Excluding the bignum arithmetic, but prominent in the list of heavily-used functions are polynomial add, subtract, quotient and remainder.

Using Gaussian elimination for determinant calculation in Macsyma takes much less time: about 11 seconds, of which 82 percent is bignum arithmetic. Using poles/residues arithmetic with Gaussian elimination on this example uses 34.4 + 6.4 (GC) = 41 seconds; given the need for division for this method, it is not surprising that poles/residues is slower.

What if we use double-precision floating point? Then the determinant is computed using minor expansion in 1.44+0.12=1.61 seconds.

Trying some other computer algebra systems on determinants provides a mixture of timing results because of the variations possible in the form of the determinant. Thus Mathematica 4.1 can “compute” the determinant in 7.8 seconds, but in a huge form that is not simplified. Combining the terms over a common denominator seems to take much longer. I stopped it at 360 seconds. Computing the partial fraction expansion form for the determinant did not seem at all feasible. There is presumably a better approach.

Using Maple 7, I tried

```
g:=proc(i,j); 1/(x+i+j-1)^j end proc;
m:=Matrix(8,8,g);
linalg[det](m);
convert(E,parfrac,x);
```

which took 69.97 seconds for the determinant calculation (returning an answer with a factored denominator), and another 23.77 seconds for converting the result to a partial-fraction form.

Since Maple is rather fast on this exact calculation, it would be interesting to see how writing a Maple program maintaining the partial-fraction form throughout the computation, would improve the result.

There is another issue by which we judge the pole/residue representation superior.

Given the two representations of a determinant as (a) expanded denominator over (say) factored denominator, or (b) sum of fractions as produced by pole/residue computation, which is more accurate to evaluate numerically?

As the size of a matrix increases, the distinction grows further, but we need only consider the 4 by 4 determinant.

Evaluate form (a) and (b) at the same double-float point -4.55. The value for (a) is 6.2d-16. The value for (b) is 321.31, accurate to 7 decimal digits. If one plots the inaccurate form (a) between -5 and -4, the space appears to be filled with huge random up-and-down jags. The form (b) simply shows a kind of steep parabolic curve with a minimum of 321 at x= -4.55.

If the two forms are rearranged so that all the coefficients are stored as double-floats, our favored form (b) still gets the right answer. Form (a) gives about -3.8d6, still quite wrong.

Either form, evaluated exactly in rational numbers produces the same answer (equivalent to 321.31007...).

As a side note, one could, from the very earliest design of Macsyma, create and add forms in residue/pole representation as part of its “general representation” for expressions. For modest-length expressions it is also essentially free. However, in a *sequence* of adds and *multiplies* requiring canonical forms, off-the-shelf Macsyma will not fare well, since an efficient program for computing and maintaining a residue/pole form for a *product*, such as we have programmed, is not built-in. Furthermore, the partial-fraction form that is built-in finds denominators irreducible (over the rationals), rather than linear over the complex algebraic numbers. (Our form *can* be produced by calling a partial-fraction expansion program after every multiply, or less often if we can determine when it is required, but this is requires zerofinding.) On the other side of

the balance, Macsyma will systematically compute with more than one variable, using multivariate factoring rather than zerofinding.

As stated earlier, profiling reveals that for our *exact* determinant benchmarks some 75 to 95 percent of the time is taken by bignum operations, which we believe can be improved (except for very low precision) through using GMP instead of native Lisp arithmetic. Some Lisps (not the one we used for these tests) use the GMP library by default. This would also substantially improve the speed of the other algorithms too.

A simple alternative for pole/residue computation that might be effective in some contexts, (at least for experimentation in computer algebra systems) is to “rename” denominators. This is an old, and probably insufficient-used tactic for reducing the complexity of tasks in computer algebra. As a simple example, by renaming  $a = 1/(x - 5)$  and  $b = 1/(x - 6)$ , the earlier equation (5) becomes

$$-28a + 126b + 7x + 7.$$

It will become important at some point to re-impose the side-relations produced from renaming, and this can be done perhaps at lower cost than by re-substitution and GCD computations. For example, knowing that  $Z = a \cdot (x - 5) - 1$  is zero, we can divide an expression to get quotient  $Q$  and remainder  $R$ :  $E = Q \cdot Z + R$ . Since  $Z = 0$ ,  $E$  may be replaced by  $R$ . Any time two expressions are multiplied, as in  $a \cdot b$ , we would have to consider whether to re-express the result as linear in  $a$  and  $b$ . In this case,  $ab = b - a$ , and the right-hand side would be preferable. Figuring out when to apply the re-substitutions in the course of a calculation is critical to keep intermediate expression swell down.

## 8 Conclusions

We have proposed, implemented, and tested a representation for univariate rational functions that can, in some circumstances, provide substantial speedup in computations. We describe two (perhaps new) approaches to multiplication of these forms. In tests of individual operations, speedups of orders of magnitude are available, most particularly in rational function addition. In a more mixed environment, in computing determinants, we see a factor of 3 or more over the best previous method in Macsyma / Lisp. We would not be surprised to see this advantage translate into other systems.

The differences we note depend heavily on the speed of the underlying arithmetic as well as the cost-saving of avoiding polynomial GCD calculations.

The benchmarks are intended to be suggestive of improvements, and do not represent the best implementations. Undoubtedly more optimized programs, and especially faster arithmetic could make our specific implementation faster.

We would expect to find similar speedups if this representation were internally coded in other computer algebra systems. We expect that these ideas are of more than theoretical interest and that practical applications can be accommodated to take advantage of this representation in spite of the limitation to univariate denominators.

In particular, what does this mean for systems such as Mathematica or Maple? We believe that writing the residue/pole code in those systems’ implementation language could provide substantial speed benefits, for places where this representation is applicable, as indicated above.

Hershey[11] describes an interesting application in which representations of various special functions are reduced, for computation, first to rational functions and then to sets of poles and residues as an aid to rapid calculation. For example,  $J_0$ , the Bessel function of the first kind of order zero

## 9 Acknowledgments

This research was supported in part by NSF grant CCR-9901933 administered through the Electronics Research Laboratory, University of California, Berkeley. The Maxima computer algebra system was invaluable in helping to debug our programs, since it provided exact partial-fraction forms for our examples.

All the programs used for testing can be obtained from the author.

## References

- [1] A-M. Bellido, “Numerical factorization of multivariate polynomials: symbolic computation of a numerical iteration function” in SNAP Proceedings (see below)
- [2] M. Bronstein and S. Lafaille, “Solutions of Linear Ordinary Differential Equations in Terms of Special Functions,” *Proc. ISSAC 2002*, ACM Press, New York, 2002. 23—28.
- [3] W. S. Brown, “On Computing with Factored Rational Expressions,” SIGSAM Bulletin 8 ACM Press (Aug. 1974).
- [4] F. Chin, “The Partial Fraction Expansion Problem and Its Inverse,” *SIAM J. Comput. Volume 6, Issue 3, (1977) 554—562*.
- [5] P. Chin, R. M. Corless, G. F. Corliss, “Optimization Strategies for the Floating-Point GCD”, *Proc. ISSAC 1998* ACM Press, New York, 1998. 228—235.
- [6] A. Diaz and E. Kaltofen. “FoxBox: A System for Manipulating Symbolic Objects in Black Box Representation,” *Proc. ISSAC 1998*, ACM Press, New York, 1998. 30—37.
- [7] Durand E., *Solutions Numeriques des Equations Algebriques*, tome 1, Masson, Paris 1960.
- [8] I. Z. Emiris, “Symbolic-numeric algebra for polynomials,” Research Report INRIA, 1997.
- [9] O. Mencer, “Rational Arithmetic Units in Computer Systems,” Ph.D dissertation, EECS, Stanford University , [stanford.edu/tr/mencer.thesis.ps.gz](http://stanford.edu/tr/mencer.thesis.ps.gz)
- [10] W.M. Gentleman and S.C. Johnson, “Analysis of Algorithms, A Case Study: Determinants of Matrices with Polynomial Entries,” *ACM Trans. on Math. Softw.* Volume 2, Issue 3 (September 1976) 232—241.
- [11] A.V. Hershey, “Approximation of Functions by Sets of Poles,” NWL Technical Report TR-2584, July, 1971. 70 pages. <http://www.dtic.mil/dtic/tr/fulltext/u2/732756.pdf>
- [12] P. Kirrinnis. “Fast Numerical Improvement of Factors of Polynomials and Partial Fractions,” *Proc. ISSAC 1998*, ACM Press, New York, 1998. 260—267.
- [13] Matlab. <http://www.mathworks.com/access/helpdesk/help/techdoc/ref/residue.shtml>
- [14] J. McNamee, “A bibliography on roots of polynomials,” (last updated to 2004) <http://www1.elsevier.com/homepage/sac/cam/mcnamee/>
- [15] V. Pan. “Numerical computation of a polynomial GCD and extensions,” Research Report 2969, INRIA, 1996.
- [16] V. Pan. “Univariate Polynomials: Nearly Optimal Algorithms for Factorization and Rootfinding”, *Proc. ISSAC 2001*, ACM Press, New York, 2001, 253—267.
- [17] SNAP Proceedings, Symbolic-Numeric Algebra for Polynomials, <http://www-sop.inria.fr/galaad/conf/1996/snap.html>.

## 10 Appendix: some Maxima programs

Note that these conversions tend to produce numbers that look “messy” to humans. For computer internal representations there is no savings in a double-precision number that can be printed accurately using only a few digits, e.g. 12.3 versus one that takes 16 digits and an exponent. A disadvantage to human consumption of the pole/residue representation is that it is more likely that the polynomial roots will have “lots of digits”.

```
/* First, some utility programs that have to do with making nearby roots "identical".
```

```
Chop take a floating point or rational number or bigfloat or complex and
rounds it to relative precision prec,
```

```
e.g. chop (3.1+4.51*i),1/10) is (9*i)/2+31/10 */
```

```
chop(z,prec):=(prec:lispr(prec),
if ?symbolp(z) then z else
  maxr( ?round(lispr(realpart(z)),prec)*prec)+%i*maxr(?round(lispr(imagpart(z)),prec)*prec))$
/* convert from maxima rational or float or bigfloat to common lisp rational */
lispr(z):= if ?numberp(z) then z else if ?atom(z) then z else if ?caar (z)=?rat then ?/(?cadr(z),?caddr
/* convert from lisp rational to maxima rational */
maxr(p) := if ?rationalp(p) then ?numerator(p)/?denominator(p) else p$
```

```
/*fppfL computes a floating point partial fraction expansion chopping roots,
and returns the result in a List*/
```

```
fppfL(p,x):=
  block ([num, den,div, res, keepfloat:false,polyfactor:false, lroots,distinctroots,qq, c],
    local(r,q,mul),
    p:rat(p), num:ratnum(p), den:ratdenom(p), div: divide(num,den),
    res: ['polypart=ratdisrep(div[1])], /* polynomial part */
    /*move leading coef of denom to num */
    num: ratdisrep(div[2]/ratcoef(den,x^hipow(den,x))),
    lroots:map(lambda([z], chop(rhs(z),10^(-6))),allroots(den)),
    mul[i]:=0, for c in lroots do mul[c]:mul[c]+1, /* multiplicity of root at location c */
    distinctroots: map(first,rest(arrayinfo(mul),2)),
    qq:rat(1),
    for c in distinctroots do qq:qq*(x-c)^mul[c], /* product of all roots^mul*/

r[c,i]:= block([f:(num/qq*(x-c)^mul[c]),j:mul[c]-i],
  if j=0 then subst(c,x,f)
  else subst (c,x, if j=0 then f else 1/j!*diff(f,x,j))),

  for c in distinctroots do for i from 1 thru mul[c] do
    res: append(res, [['residue=r[c,i], 'pole=c, 'pow=i]]),
  return(res))$
```

```
/* this program is similar to the one above but returns a partial fraction form as
a Maxima expression, and avoids ratsimp or gcds after the initial division of num/denom.
*/
```

```
fppf(p,x):= /*floating point partial fraction expansion chopping roots */
  block ([num, den,div, res, keepfloat:false,polyfactor:false, lroots,distinctroots, c],
    local(r,mul,qq),
    p:rat(p), num:ratnum(p), den:ratdenom(p), div: divide(num,den),
    res: ratdisrep(div[1]), /* polynomial part */
    /*move leading coef of denom to num */
    num: ratdisrep(div[2]/ratcoef(den,x^hipow(den,x))),
    lroots:map(lambda([z], chop(rhs(z),10^(-3))),allroots(den)),
```



```

mul[i]:=0, for c in lroots do mul[c]:mul[c]+1, /* multiplicity of root at location c */
distinctroots: map(first,rest(arrayinfo(mul),2)),

qq[x]:=1,
for c in distinctroots do
  for c2 in distinctroots do if c#c2 then qq[c]: qq[c]*(x-c2)^mul[c2],
/* product of all but one roots^mul*/
r[c,i]:= block([f:(num/qq[c]),j:mul[c]-i],
  if j=0 then subst(c,x,f)
  else subst (c,x, if j=0 then f else 1/j!*diff(f,x,j))),

for c in distinctroots do for i from 1 thru mul[c] do
  res: res+r[c,i]/(x-c)^i,
return(res))$

/* try to make the denoms look monic.. start with a/(b*x+c). return (a/b) / (x+ c/b) in floats.
*/
mon(p,x):=
block([r:rat(p),b],
  b:bothcoef(ratdenom(r),x),
  rectform( ev(ratdisrep(ratnumber(r)/b[1]),numer)) /
(x+ rectform(ev(ratdisrep( b[2]/b[1]),numer)))

better form

fppf(p,x):= /*floating point partial fraction expansion chopping roots */
block ([num, den,div, res, keepfloat:false,polyfactor:false, lroots,distinctroots, c],
  local(r,mul,qq),
  p:rat(p), num:ratnumber(p), den:ratdenom(p), div: divide(num,den),
  res: ratdisrep(div[1]), /* polynomial part */
/*move leading coef of denom to num */
num: ratdisrep(div[2]/ratcoef(den,x^hipow(den,x))),
  lroots:map(lambda([z], chop(rhs(z),10^(-3))),allroots(den)),

  mul[i]:=0, for c in lroots do mul[c]:mul[c]+1, /* multiplicity of root at location c */
  distinctroots: map(first,rest(arrayinfo(mul),2)),
/* distinctroots:map('float,distinctroots),*/

  qq[x]:=1,
  for c in distinctroots do
    for c2 in distinctroots do if c#c2 then qq[c]: qq[c]*(x-c2)^mul[c2],
/* product of all but one roots^mul*/
r[c,i]:= block([f:(num/qq[c]),j:mul[c]-i],
  if j=0 then subst(c,x,f)
  else subst (c,x, if j=0 then f else 1/j!*diff(f,x,j))),

  for c in distinctroots do for i from 1 thru mul[c] do
    res: res+ev(rectform(expand(r[c,i]),bfloat)/(x-ev(c,numer))^i,
return(res))$

```

A Pade approximation to sine, expanded around zero

$$\frac{479249 x^7 - 52785432 x^5 + 1640635920 x^3 - 11511339840 x}{18361 x^6 + 3177720 x^4 + 277920720 x^2 + 11511339840}$$