

FAST AND RIGOROUS ARBITRARY-PRECISION COMPUTATION OF GAUSS-LEGENDRE QUADRATURE NODES AND WEIGHTS

FREDRIK JOHANSSON* AND MARC MEZZAROBBA†

Abstract. We describe a strategy for rigorous arbitrary-precision evaluation of Legendre polynomials on the unit interval and its application in the generation of Gauss-Legendre quadrature rules. Our focus is on making the evaluation practical for a wide range of realistic parameters, corresponding to the requirements of numerical integration to an accuracy of about 100 to 100 000 bits. Our algorithm combines the summation by rectangular splitting of several types of expansions in terms of hypergeometric series with a fixed-point implementation of Bonnet’s three-term recurrence relation. We then compute rigorous enclosures of the Gauss-Legendre nodes and weights using the interval Newton method. We provide rigorous error bounds for all steps of the algorithm. The approach is validated by an implementation in the Arb library, which achieves order-of-magnitude speedups over previous code for computing Gauss-Legendre rules with simultaneous high degree and precision.

Key words. Legendre polynomials, Gauss-Legendre quadrature, arbitrary-precision arithmetic, interval arithmetic

AMS subject classifications. 65Y99, 65G99, 33C45

1. Introduction. The Legendre polynomials $P_n(x)$ are the sequence of orthogonal polynomials with respect to the unit weight on the interval $(-1, 1)$, normalized so that $P_n(1) = 1$. Like other classical orthogonal polynomials, Legendre polynomials satisfy a three-term recurrence, in this case the relation

$$(1) \quad (n+1)P_{n+1}(x) - (2n+1)xP_n(x) + nP_{n-1}(x) = 0,$$

also known as Bonnet’s formula, and a second order differential equation, here

$$(2) \quad (1-x^2)P_n''(x) - 2xP_n'(x) + n(n+1)P_n(x) = 0.$$

The definition implies that P_n has n roots all located in $(-1, 1)$. Perhaps the most important application of Legendre polynomials is the Gauss-Legendre quadrature rule

$$(3) \quad \int_{-1}^1 f(x)dx \approx \sum_{i=0}^{n-1} w_i f(x_i), \quad w_i = \frac{2}{(1-x_i^2)[P_n'(x_i)]^2},$$

where the *nodes* x_i are the roots of P_n . The quantity w_i is called the *weight* associated with the node x_i .

For some applications in computer algebra, number theory, mathematical physics, and experimental mathematics, it is necessary to compute integrals to an accuracy of hundreds of digits, and occasionally even tens of thousands of digits [1]. It is well known that the Gauss-Legendre formula (3) is nearly optimal among all quadrature rules for integrands that are well-approximated by polynomials, in particular for analytic f without singularities close to the path of integration [17, 27]. As a special case, the n -point rule (3) is exact when f is a polynomial of degree at most $2n - 1$. In general, the error in (3) can be bounded in terms of $\sup_{x \in (-1, 1)} |f^{(2n)}(x)|$, or if f is analytic on a domain D containing $(-1, 1)$, in terms of $\sup_{z \in D} |f(z)|$ and the

*INRIA – LFANT, CNRS – IMB – UMR 5251, Université de Bordeaux, 33400 Talence, France (fredrik.johansson@gmail.com)

†Sorbonne Université, CNRS, Laboratoire d’Informatique de Paris 6, LIP6, F-75005 Paris, France (marc@mezzarobba.net)

distance from the boundary of D to $(-1, 1)$. For integration of smooth functions without endpoint singularities, Gauss-Legendre quadrature achieves an accuracy of p bits with $n = \mathcal{O}(p)$ evaluation points. Even when the conditions are not ideal for using (3) directly, rapid convergence is often possible by combining (3) with adaptive subdivision of the integration path [22].

The Gauss-Legendre scheme has the drawback that the quadrature nodes and weights are somewhat inconvenient to compute. Indeed, P_n becomes highly oscillatory for large n and hence presents difficulties for naive root-finding and polynomial evaluation methods. The classical Golub-Welsch algorithm avoids accuracy problems by formulating the task of computing the nodes as finding the eigenvalues of a tridiagonal matrix [9], but this approach is still too slow to be practical for large n .

In the last decade, several authors have contributed to the development of asymptotic methods that permit computing any individual node and weight for arbitrarily large n in $\mathcal{O}(1)$ time, culminating in the 2014 work by Bogaert [4, 11, 3]. For a review of this progress, see Townsend [26]. Of course, the “ $\mathcal{O}(1)$ ” bound assumes that a fixed level of precision is used. In the prevailing literature this generally means 53-bit IEEE 754 floating-point arithmetic. As mentioned earlier, certain applications require considering $p \sim n$ where p is the precision in bits, which potentially can be in the thousands. In addition, the available $\mathcal{O}(1)$ implementations rely in part on heuristic error estimates without rigorously proved bounds.

The literature on arbitrary precision or rigorous evaluation is comparatively limited. Petras [21] gave explicit bounds for the error $|x_k^{(i)} - x_k|$ when the roots x_k of the Legendre polynomial P_n are approximated using Newton iteration

$$(4) \quad x_k^{(i+1)} = x_k^{(i)} - \frac{P_n(x_k^{(i)})}{P_n'(x_k^{(i)})}$$

provided that the initial values $x_k^{(0)}$ are computed by a certain asymptotic formula. However, Petras did not address the numerical evaluation of $P_n(x)$. Fousse [8] discussed the rigorous implementation of Gauss-Legendre quadrature using generic polynomial root isolation methods together with interval Newton iteration for root refinement, but did not study fast methods for large n . Code for high-precision Gauss-Legendre quadrature rules can also be found in packages such as Pari/GP [25] and ARPREC [2], but without error analysis and without special techniques for large n .

If we assume that the precision p varies, then it is clear that any node and weight can be computed to p -bit accuracy in $\tilde{\mathcal{O}}(np)$ time¹, by performing $\mathcal{O}(\log p)$ Newton iterations (4) from an appropriate initial value. As a consequence, the full set of nodes and weights for the degree- n quadrature rule can be computed in $\tilde{\mathcal{O}}(n^2p)$ time. For numerical integration of analytic functions where we typically have $p \sim n$, a better (indeed, optimal) estimate than the classical $\tilde{\mathcal{O}}(n^3)$ bound is possible.

THEOREM 1. *If $p \sim n$, then the Gauss-Legendre nodes and weights of degree n can be computed to p -bit accuracy in $\tilde{\mathcal{O}}(n^2)$ (equivalently, $\tilde{\mathcal{O}}(p^2)$) bit operations.*

Proof sketch. Using the formulas in [21], we can compute good initial values for Newton iteration in $\tilde{\mathcal{O}}(n)$ bit operations. The Newton iteration can be performed for all roots simultaneously using fast multipoint evaluation, which costs $\tilde{\mathcal{O}}(np)$ bit operations. Fast multipoint evaluation is numerically unstable and generically loses

¹We write $\tilde{\mathcal{O}}(x)$ for $\mathcal{O}(x \log(x)^{\mathcal{O}(1)})$.

$\mathcal{O}(n)$ bits of accuracy, but we can compensate for this loss by using $\mathcal{O}(n)$ guard bits [16]. Since $p \sim n$ by assumption, this does not change the complexity bound. \square

Completing the details of the proof is a technical exercise. Despite being elegant in theory, the algorithm behind [Theorem 1](#) has high overhead in practice, in large part due to the need to work with greatly increased precision. Working with expanded polynomials and processing all roots simultaneously also results in high memory usage and makes parallelization difficult. See [subsection 4.4](#) below for an alternative approach—also unlikely to be practical except for extremely large p —that leads to a similar complexity bound. We can achieve a slightly worse but still sub-cubic complexity of $\tilde{\mathcal{O}}(n^{5/2})$ by employing fast multipoint evaluation in a completely different way to compute P_n values in isolation, but unfortunately that algorithm also has high overhead [12].

In this work, we study fast and practical computation for variable n and precision p with rigorous error bounds. Our main contribution is to give a complete evaluation strategy for Legendre polynomials on $[-1, 1]$ in ball arithmetic [28, 15]. Computing the Legendre polynomial roots, then, is a relatively simple application of the results in [21] together with the interval Newton method [18].

Our algorithm for evaluating Legendre polynomials switches between different methods. In [section 3](#), we prove practical error bounds for the three-term recurrence (1), which can be efficiently implemented in fixed-point arithmetic. This method is ideal for n and p up to a few hundred. For larger n or p , we use a fast method for evaluation of hypergeometric series. [Section 4](#) discusses the hypergeometric series expansions that are preferable for different inputs and precision (including a well-known asymptotic expansion for large n), and their efficient evaluation. In [section 5](#), we propose a strategy to select the best formula for any combination of n, p, x .

[Section 6](#) presents benchmark results that compare the performance of our algorithm to some previous implementations as well as the asymptotically fast algorithm in [Theorem 1](#).

For generating Gauss-Legendre quadrature rules with $n \sim p$, our algorithm has an asymptotic complexity of $\tilde{\mathcal{O}}(n^3)$ like classical methods. However, our algorithm has much lower overhead, and for parameters $p, n < 10^5$ which are most relevant to applications, the observed running time is effectively subcubic. Furthermore, if $p = \mathcal{O}(1)$, the complexity reduces to $\tilde{\mathcal{O}}(n)$ as in the machine-precision implementations by Bogaert and others.

Finally, [section 7](#) reviews the viability of Gauss-Legendre quadrature compared to other methods for extremely high precision integration.

Our code for evaluating Legendre polynomials and computing Gauss-Legendre nodes and weights is freely available as part of the Arb library [15].

2. General strategy. We work in the framework on midpoint-radius interval arithmetic, also called ball arithmetic. In general, given an integer n and a ball $x = [m \pm r] = [m - r, m + r]$, we want to evaluate $P_n(x)$ at x , yielding an enclosure $y = [m' \pm r']$ such that $P_n(\xi) \in y$ holds for all $\xi \in x$.

We restrict our attention to real $x \in [-1, 1]$, which is the most interesting part of the domain for applications. Since $P_n(-x) = (-1)^n P_n(x)$, we can further restrict to $0 \leq x \leq 1$. Some authors suggest working with $P_n(\cos(\theta))$ instead of $P_n(x)$ directly to improve numerical stability for x close to 1, but this is not necessary in arbitrary-precision arithmetic since a slight precision increase (of the order of $\mathcal{O}(\log n)$ bits) works as well.

We note that, for rigorous evaluation of $P_n(z)$ with complex z as well as Legendre

functions of non-integer order n , generic methods for the hypergeometric ${}_2F_1$ function are applicable if n is not extremely large; see [14]. Real $|x| > 1$ can also be handled easily using naive methods.

In view of the use of Newton's method to compute the roots, we also need to evaluate the derivative $P'_n(x)$, typically at the same time as $P_n(x)$ itself. A simple option is to deduce $P'_n(x)$ from $P_n(x)$ and $P_{n-1}(x)$ using

$$(5) \quad (x^2 - 1)P'_n(x) = n(xP_n(x) - P_{n-1}(x)).$$

When x is close to 1, though, this formula involves a cancellation of about $|\log_2(1-x)|$ bits in the subtraction, followed by a division by $x^2 - 1$. Therefore, a direct evaluation of $P'_n(x)$ may be preferable to reduce the working precision.

Our evaluation algorithms rely on ball arithmetic internally to propagate the error bounds up to the final result. Since some methods would produce unsatisfactorily large enclosures when executed on input balls $x = [m \pm r]$ of radius $r > 0$, we evaluate $P_n(m)$ (with higher internal precision if necessary) and use a first-order bound

$$\max_{\xi \in x} |P_n(\xi) - P_n(m)| \leq r \max_{\xi \in x} |P'_n(\xi)|$$

to separately bound the propagated error. Similarly, we use a bound for P''_n to compute a reasonably tight enclosure for $P'_n([m \pm r])$. Suitable bounds are given in [Proposition 3](#) below.

LEMMA 2. Denote by $[z^n]f(z)$ the coefficient of index n in a power series $f(z)$, and write $f(z) \ll_z \hat{f}(z)$ if $|[z^n]f(z)| \leq [z^n]\hat{f}(z)$ for all n . If f, g, \hat{f}, \hat{g} are such that $f(z) \ll_z \hat{f}(z)$ and $g(z) \ll_z \hat{g}(z)$, then we also have $\int_0^z f \ll_z \int_0^z \hat{f}$ and $f(z)g(z) \ll_z \hat{f}(z)\hat{g}(z)$.

PROPOSITION 3. The following bounds hold for $-1 \leq x \leq 1$:

$$(6) \quad |P'_n(x)| \leq \min \left(\frac{2^{3/2}}{\sqrt{\pi}} \frac{\sqrt{n}}{(1-x^2)^{3/4}}, \frac{n(n+1)}{2} \right),$$

$$(7) \quad |P''_n(x)| \leq \min \left(\frac{2^{5/2}}{\sqrt{\pi}} \frac{n^{3/2}}{(1-x^2)^{5/4}}, \frac{(n-1)n(n+1)(n+2)}{8} \right).$$

Proof. It is classical that Legendre polynomials are given by the generating series

$$(8) \quad F(x, z) = \sum_{n=0}^{\infty} P_n(x)z^n = \frac{1}{\sqrt{1-2xz+z^2}}.$$

Differentiation with respect to x yields

$$\sum_{n=0}^{\infty} P'_n(x)z^n = \frac{zF(x, z)}{1-2xz+z^2}, \quad \sum_{n=0}^{\infty} P''_n(x)z^n = \frac{3z^2F(x, z)}{(1-2xz+z^2)^2}.$$

Set $\theta = \arccos x$, so that the roots of $1-2xz+z^2$ are $e^{\pm i\theta}$. Then, in the notation of [Lemma 2](#), we have the bound

$$\frac{1}{1-2xz+z^2} = \frac{1}{2i \sin \theta} \left(\frac{1}{z-e^{i\theta}} - \frac{1}{z-e^{-i\theta}} \right) = \sum_{n=0}^{\infty} \frac{\sin((n+1)\theta)}{\sin(\theta)} z^n \ll_z \frac{\sin(\theta)^{-1}}{1-z}.$$

In addition, Bernstein's inequality for the Legendre polynomials (see, e.g., [6]) combined with the logarithmic convexity of the Gamma function yields

$$|P_n(x)| \leq \frac{\sqrt{2}}{\sqrt{\pi}} \frac{1}{\sqrt{\sin \theta}} \frac{1}{\sqrt{n+1/2}} \leq \frac{\sqrt{2}}{\sqrt{\pi}} \frac{1}{\sqrt{\sin \theta}} \frac{\Gamma(n+1/2)}{\Gamma(n+1)},$$

and hence

$$F(x, z) \ll_z \frac{\sqrt{2}}{\sqrt{\sin \theta}} \sum_{n=0}^{\infty} \frac{1}{\sqrt{\pi}} \frac{\Gamma(n+1/2)}{\Gamma(n+1)} z^n = \frac{\sqrt{2}}{\sqrt{\sin \theta}} \frac{1}{\sqrt{1-z}}.$$

By Lemma 2, these bounds combine into

$$\frac{dF}{dx} \ll_z \frac{\sqrt{2}}{\sin(\theta)^{3/2}} \frac{1}{(1-z)^{3/2}}, \quad \frac{d^2F}{dx^2} \ll_z \frac{\sqrt{2}}{\sin(\theta)^{5/2}} \frac{3z^2}{(1-z)^{5/2}}.$$

Since $[z^n](1-z)^{-k-1/2} = \Gamma(n+1/2)/(\Gamma(k+1/2)\Gamma(n-k+1))$ and using again the logarithmic convexity of Γ , we conclude that

$$|P'_n(x)| \leq \frac{\sqrt{2}}{\sin(\theta)^{3/2}} \frac{2}{\sqrt{\pi}} \frac{\Gamma(n+1/2)}{\Gamma(n)} \leq \frac{2^{3/2}}{\sqrt{\pi}} \frac{\sqrt{n}}{\sin(\theta)^{3/2}}, \quad |P''_n(x)| \leq \frac{2^{5/2}}{\sqrt{\pi}} \frac{n^{3/2}}{(1-x^2)^{5/4}}.$$

The result follows since all derivatives of Legendre polynomials reach their maximum at $x = 1$ (or by using the bounds $(z - e^{\pm i\theta})^{-1}$, $F(x, z) \ll_z (1-z)^{-1}$ and Lemma 2). \square

Remark 4. By the same reasoning, the inequality

$$|P_n^{(k)}(x)| \leq \frac{2^{k+1/2}}{\sqrt{\pi}} \frac{n^{k-1/2}}{(1-x^2)^{(2n+1)/4}}$$

actually holds for all k . Unfortunately, it seems to overestimate $|P^{(k)}|$ by a factor about 2^k .

Based on these reductions, we assume from now on that x is a floating-point number with $0 \leq x \leq 1$. Our main algorithm for evaluating P_n at x combines the following methods:

- the iterative computation of $P_n(x)$ via the three-term recurrence (1),
- an asymptotic expansion of $P_n(x)$ as $n \rightarrow \infty$,
- the usual expanded expression of P_n in the monomial basis,
- the analogous terminating expansion at 1.

All three expansions can be written as hypergeometric series, i.e., sums of the form $\sum_k c_k \xi^k$ where c_k/c_{k-1} is a rational function of k .

The constraints and heuristics used to select between these methods are described in detail below. Roughly speaking, the three-term recurrence is used for small index n and precision p , when x is not too close to 1; the asymptotic series when n is large enough, again with x not too close to 1; the expansion at 0 for large p unless x is close to 1; and finally the expansion at 1 in the remaining cases when x is close to 1.

At p -bit precision, it is sufficient to choose algorithms and internal evaluation parameters to target an absolute error of about 2^{-p} , or more precisely $2^{-p - \mathcal{O}(\log n)}$ for a relative error of about 2^{-p} measured with respect to monotone envelopes for $P_n(x)$ and $P'_n(x)$ as in [4]. The relative error of a computed ball for $P_n(x)$ where x is near a zero x_k can be arbitrarily large, but the relative error of $P'_n(x)$ near x_k will be small, which is sufficient for Newton iteration. Since the output consists of a ball, we also have the option of catching a result with large relative error and repeating the evaluation with a higher precision as needed.

3. Basecase recurrence. For small n , a straightforward way to compute $P_n(x)$ is to apply the three-term recurrence (1), starting from $P_0(x) = 1$ and $P_1(x) = x$. Computing $P_n(x)$ by this method takes about $(\mathcal{M}(t) + \mathcal{O}(t))n$ bit operations, where t is the working precision and $\mathcal{M}(t)$ denotes the cost of t -bit multiplication. It is thus attractive for small n and t , especially when both $P_n(x)$ and $P'_n(x)$ are needed, since we can get $P_{n-1}(x)$ at no additional cost.

Fix $x \in [-1, 1]$, and let $p_n = P_n(x)$. Bonnet's formula (1) gives

$$(9) \quad p_{n+1} = \frac{1}{n+1} \left((2n+1)xp_n - np_{n-1} \right), \quad n \geq 0.$$

In a direct implementation of this recurrence in ball arithmetic, the width of the enclosures would roughly double at every iteration, requiring to increase the internal working precision by $\mathcal{O}(n)$ bits. We avoid this issue by performing a pen-and-paper round-off error analysis of the evaluation that yields a less pessimistic bound on the accumulated error. Additionally, the static error bound allows us to implement the recurrence in fixed-point arithmetic, avoiding the overhead of floating-point and interval operations.

Suppose $x = \hat{x}2^{-t}$ with $\hat{x} \in \mathbb{Z}$ is a given fixed-point number. Let $\lceil u \rceil$ denote the integer truncation of a real number u (any other rounding function would do). The integer sequence (\hat{p}_n) defined by

$$(10) \quad \hat{p}_0 = 2^t, \quad \hat{p}_1 = \hat{x}, \quad \hat{p}_{n+1} = \left\lceil \frac{1}{n+1} \left((2n+1)\lceil \hat{x}\hat{p}_n2^{-t} \rceil - n\hat{p}_{n-1} \right) \right\rceil$$

is easy to compute using only integer arithmetic, and \hat{p}_n2^{-t} is an approximation of p_n . Algorithm 1 provides a complete C implementation using GMP [10]. As a small optimization, we delay the division by $n+1$ until we have accumulated a denominator of the size of a machine word.

To bound the difference $|\hat{p}_n2^{-t} - p_n|$, we analyze the effect on the result of a small perturbation in each iteration of (9). The bound is based on a classical linearity argument (compare, e.g., [29]) combined with majorant series techniques.

PROPOSITION 5. *Suppose that a sequence $(\tilde{p}_n)_{n \geq -1}$ satisfies $\tilde{p}_0 = 1$ and*

$$(11) \quad \tilde{p}_{n+1} = \frac{1}{n+1} \left((2n+1)x\tilde{p}_n - n\tilde{p}_{n-1} \right) + \varepsilon_n, \quad n \geq 0.$$

for arbitrary real numbers ε_n with $|\varepsilon_n| \leq \bar{\varepsilon}$ for all n . Then we have

$$|\tilde{p}_n - P_n(x)| \leq \frac{(n+1)(n+2)}{4} \bar{\varepsilon}$$

for all $n \geq 0$.

Proof. Let $\delta_n = \tilde{p}_n - p_n$ and $\eta_n = (n+1)\varepsilon_n$. Subtracting (9) from (11) gives

$$(12) \quad (n+1)\delta_{n+1} = (2n+1)x\delta_n - n\delta_{n-1} + \eta_n,$$

with $\delta_0 = 0$. Consider the formal generating series $\delta(z) = \sum_{n \geq 0} \delta_n z^n$ and $\eta(z) = \sum_{n \geq 0} \eta_n z^n$. Noting that (12) holds for all $n \in \mathbb{Z}$ if the sequences (δ_n) and (η_n) are extended by 0 for $n < 0$ and using the relations

$$\sum_{n=-\infty}^{\infty} f_{n-1}z^n = z \sum_{n=-\infty}^{\infty} f_n z^n, \quad \sum_{n=-\infty}^{\infty} n f_n z^n = z \frac{d}{dz} \sum_{n=-\infty}^{\infty} f_n z^n,$$

Algorithm 1 Evaluation of Legendre polynomials in GMP fixed-point arithmetic

Input: An integer x and $t \geq 0$ such that $|2^{-t}x| \leq 1$, and $n \geq 1$

Output: p, q such that $|2^{-t}p - P_{n-1}(2^{-t}x)|, |2^{-t}q - P_n(2^{-t}x)| \leq (0.75(n+1)(n+2)+1)2^{-t}$

```
void legendre(mpz_t p, mpz_t q, int n, const mpz_t x, int t) {
    mpz_t tmp; int k; mpz_init(tmp);
    mp_limb_t denlo, den = 1;
    mpz_set_ui(p, 1);
    mpz_mul_2exp(p, p, t);
    mpz_set(q, x);
    for (k = 1; k < n; k++) {
        mpz_mul(tmp, q, x);
        mpz_tdiv_q_2exp(tmp, tmp, t);
        mpz_mul_si(p, p, -k*k);
        mpz_addmul_ui(p, tmp, 2*k+1);
        mpz_swap(p, q);
        if (mpn_mul_1(&denlo, &den, 1, k+1)) {
            mpz_tdiv_q_ui(p, p, den);
            mpz_tdiv_q_ui(q, q, den);
            den = k+1;
        } else den = denlo;
    }
    mpz_tdiv_q_ui(p, p, den/n);
    mpz_tdiv_q_ui(q, q, den);
    mpz_clear(tmp);
}
```

we see that (12) translates into

$$(1 - 2xz + z^2)z \frac{d}{dz} \delta(z) = z(x - z)\delta(z) + z\eta(z).$$

The solution of this differential equation with $\delta(0) = 0$ reads, cf. (8),

$$\delta(z) = p(z) \int_0^z \eta(w) p(w) dw, \quad p(z) = \sum_{n=0}^{\infty} p_n z^n = F(x, z) = \frac{1}{\sqrt{1 - 2xz + z^2}}.$$

This is an exact expression of the “global” error δ in terms of the “local” errors ε_n . Since $|p_n| = |P_n(x)| \leq 1$ and $|\eta_n| \leq (n+1)\bar{\varepsilon}$, it follows by Lemma 2 that

$$|\delta_n| = \left| [z^n] \left(p(z) \int_0^z \eta(w) p(w) dw \right) \right| \leq [z^n] \left(\frac{1}{1-z} \int_0^z \frac{\bar{\varepsilon}}{(1-w)^2} \frac{1}{1-w} dw \right)$$

and therefore

$$|\delta_n| \leq [z^n] \left(\frac{1}{2} \frac{\bar{\varepsilon}}{(1-z)^3} \right) = \frac{(n+1)(n+2)}{4} \bar{\varepsilon}. \quad \square$$

COROLLARY 6. *Suppose that $x = \hat{x}2^{-t}$ for some $t \geq 0$ and $\hat{x} \in \mathbb{Z}$. The sequence $(\hat{p}_n)_{n \geq 0}$ defined by (10) satisfies*

$$(13) \quad |\hat{p}_n 2^{-t} - p_n| \leq 0.75(n+1)(n+1)2^{-t}, \quad n \geq 0.$$

Furthermore, the quantities p, q returned by Algorithm 1 are such that

$$(14) \quad |p - 2^t P_{n-1}(x)|, |q - 2^t P_n(x)| \leq 0.75(n+1)(n+2) + 1.$$

Proof. We can write

$$\hat{p}_{n+1} = \frac{1}{n+1} \left((2n+1)(\hat{x} \hat{p}_n 2^{-t} + \alpha_n) - n \hat{p}_{n-1} \right) + \beta_n$$

for some α_n, β_n of absolute value at most one, and hence

$$\hat{p}_{n+1} = \frac{1}{n+1} \left((2n+1) \hat{x} \hat{p}_n 2^{-t} - n \hat{p}_{n-1} \right) + \varepsilon_n, \quad \varepsilon_n = \frac{2n+1}{n+1} \alpha_n + \beta_n.$$

where $|\varepsilon_n| \leq 3$. [Proposition 5](#) applied to $\tilde{p}_n = \hat{p}_n 2^{-t}$ then provides the bound [\(13\)](#).

Turning to [Algorithm 1](#), let $\mathbf{p}_0, \mathbf{q}_0, \mathbf{d}_0$ denote the values of the variables $\mathbf{p}, \mathbf{q}, \mathbf{den}$ before the loop, and $\mathbf{p}_k, \mathbf{q}_k, \mathbf{d}_k$ their values at the end of iteration k . Consider the sequence $\tilde{p}_k = 2^{-t} \mathbf{q}_{k-1} / \mathbf{d}_{k-1}$, $k \geq 1$, extended by $\tilde{p}_0 = 1$ and an arbitrary \tilde{p}_{-1} . For all $k \geq 1$, depending whether the conditional branch is taken, we have one of the systems of equations

$$(15) \quad \mathbf{p}_k = \mathbf{q}_{k-1}, \quad \mathbf{q}_k = (2k+1) \lceil \hat{x} \mathbf{q}_{k-1} 2^{-t} \rceil - k^2 \mathbf{p}_{k-1}, \quad \mathbf{d}_k = (k+1) \mathbf{d}_{k-1}$$

$$(16) \quad \mathbf{p}_k = \left\lfloor \frac{\mathbf{q}_{k-1}}{\mathbf{d}_{k-1}} \right\rfloor, \quad \mathbf{q}_k = \left\lfloor \frac{(2k+1) \lceil \hat{x} \mathbf{q}_{k-1} 2^{-t} \rceil - k^2 \mathbf{p}_{k-1}}{\mathbf{d}_{k-1}} \right\rfloor, \quad \mathbf{d}_k = k+1.$$

In both cases, we can write

$$\frac{\mathbf{p}_k}{\mathbf{d}_k} = \frac{\mathbf{q}_{k-1}}{(k+1) \mathbf{d}_{k-1}} + \frac{\alpha_k}{k+1}, \quad \frac{\mathbf{q}_k}{\mathbf{d}_k} = \frac{(2k+1)(x \mathbf{q}_{k-1} + \beta_k) - k^2 \mathbf{p}_{k-1}}{(k+1) \mathbf{d}_{k-1}} + \frac{\gamma_k}{k+1}$$

with $|\alpha_k|, |\beta_k|, |\gamma_k| \leq 1$. The first equation implies $2^{-t} k \mathbf{p}_{k-1} / \mathbf{d}_{k-1} = \tilde{p}_{k-1} + \alpha_{k-1} 2^{-t}$ for $k \geq 2$. Since the latter equality also holds for $k = 1$ with $\alpha_0 = 0$, we can substitute it in the second equation, yielding

$$\tilde{p}_{k+1} = \frac{(2k+1)x\tilde{p}_k - k\tilde{p}_{k-1}}{k+1} + 2^{-t} \left(\frac{-k\alpha_{k-1}}{k+1} + \frac{(2k+1)\beta_k}{(k+1)\mathbf{d}_{k-1}} + \frac{\gamma_k}{k+1} \right).$$

This relation holds for $k \geq 1$, and we extend it to $k = 0$ by setting $\beta_0 = \gamma_0 = 0$. Thus, \tilde{p}_k also satisfies [\(11\)](#) with $|\varepsilon_n| \leq 3 \cdot 2^{-t}$, and [Proposition 5](#) applies. The final values of \mathbf{q} and \mathbf{p} are respectively $\lceil 2^t \tilde{p}_n \rceil$ and $\lceil n \mathbf{p}_{n-1} / \mathbf{d}_{n-1} \rceil = \lceil 2^t \tilde{p}_{n-1} + \alpha_{n-1} \rceil$, whence the bound [\(14\)](#). \square

We do not use asymptotically faster evaluation techniques for large n in combination with this recurrence, since the series expansions to be presented next perform very well in this case.

4. Series expansions. For large n or p , we employ series expansions of $P_n(x)$ with respect to either n or x rather than the algorithm from the previous section. The coefficients of the series are also computed by recurrence, but fewer than n terms will typically be required. Additionally, the expansions are of a form suitable for fast evaluation by rectangular splitting. Let us first review the various series expansions that we are using (an asymptotic expansion as $n \rightarrow \infty$, series expansions at $x = 0$ and $x = 1$), and then discuss their efficient evaluation.

4.1. Asymptotic series. For fixed $|x| < 1$ or equivalently $x = \cos(\theta)$ with $0 < \theta < \pi$, an asymptotic expansion for $P_n(x)$ as $n \rightarrow \infty$ can be given as [\[4, Eq. 3.4\]](#)

$$(17) \quad P_n(\cos(\theta)) = \left(\frac{2}{\pi \sin(\theta)} \right)^{1/2} \sum_{k=0}^{K-1} C_{n,k} \frac{\cos(\alpha_{n,k}(\theta))}{\sin^k(\theta)} + \xi_{n,K}(\theta)$$

where

$$(18) \quad C_{n,k} = \frac{[\Gamma(k+1/2)]^2 \Gamma(n+1)}{\pi 2^k \Gamma(n+k+3/2) \Gamma(k+1)},$$

$$(19) \quad \alpha_{n,k}(\theta) = (n+k+1/2)\theta - (k+1/2)\pi/2,$$

and the error term satisfies

$$(20) \quad |\xi_{n,K}(\theta)| < 2 \left(\frac{2}{\pi \sin(\theta)} \right)^{1/2} \frac{C_{n,K}}{\sin^K(\theta)}.$$

The coefficients $C_{n,k}$ are a hypergeometric sequence with

$$(21) \quad \frac{C_{n,k}}{C_{n,k-1}} = \frac{(2k-1)^2}{k(2n+2k+1)}, \quad C_{n,0} = \frac{1}{\sqrt{\pi}} \frac{4^n}{(n+\frac{1}{2}) \binom{2n}{n}}.$$

To evaluate the error bound, we can use the following inequality for $n, k \geq 1$:

$$(22) \quad C_{n,k} \leq \frac{1}{\pi n^{1/2}} \frac{k!n!}{2^k(n+k)!} \leq \frac{1}{\pi n^{1/2}} \frac{k!}{(2n)^k}.$$

The asymptotic expansion is in fact a convergent series when $2 \sin(\theta) > 1$, which allows evaluating $P_n(x)$ to unbounded accuracy for fixed n when $\frac{1}{6}\pi < \theta < \frac{5}{6}\pi$. The particular form (17) must be used for this purpose; there is a slightly different version of the expansion which is asymptotic to $P_n(x)$ (for fixed K when $n \rightarrow \infty$) but paradoxically converges to $2P_n(x)$ (for fixed n when $K \rightarrow \infty$); see [19] and [20].

We can restate (17) as a hypergeometric series by working with complex numbers. Letting $\omega = 1 - (x/y)i$, with $x = \cos(\theta)$ and $y = \sin(\theta)$ as usual, we have

$$(23) \quad P_n(x) = \sqrt{\pi y} \operatorname{Re} \left[(1-i)(x+yi)^{n+1/2} \sum_{k=0}^{K-1} C_{n,k} \omega^k \right] + \xi_{n,K}(\theta).$$

This eliminates the explicit trigonometric functions and permits using [Algorithm 2](#) below to evaluate the series.

The evaluation of (23) in ball arithmetic is numerically stable, and we therefore only need to add a few guard bits to the working precision.

4.2. Expansion at zero. If $n = 2d$ is even, the expansion of $P_n(x)$ in the monomial basis reads

$$(24) \quad \begin{aligned} P_{2d}(x) &= (-1)^d \sum_{k=0}^d \frac{(-1)^k}{2^n} \binom{n}{d-k} \binom{n+2k}{n} x^{2k} \\ &= \frac{(-1)^d}{2^{2d}} \binom{2d}{d} \sum_{k=0}^d A_{-1}(d,k) (-x^2)^k, \end{aligned}$$

and if $n = 2d + 1$ is odd, we have

$$(25) \quad \begin{aligned} P_{2d+1}(x) &= (-1)^d x \sum_{k=0}^d \frac{(-1)^k}{2^n} \binom{n}{d-k} \binom{n+2k+1}{n} x^{2k} \\ &= \frac{(-1)^d (d+1)}{2^{2d+1}} \binom{2d+2}{d+1} x \sum_{k=0}^d A_{+1}(d,k) (-x^2)^k \end{aligned}$$

where the hypergeometric sequences $A_{\pm 1}$ can be defined by $A_{\pm 1}(d, 0) = 1$ and

$$(26) \quad \frac{A_{\sigma}(d, k)}{A_{\sigma}(d, k-1)} = \frac{(d-k+1)(2d+2k+\sigma)}{k(2k+\sigma)}, \quad \sigma \in \{-1, +1\}.$$

At very high precision, we evaluate the full polynomials, where (24) and (25) have the advantage compared to other expansions of requiring only $n/2$ terms due to the odd-even form. At lower precision p , the high order terms will be smaller than 2^{-p} when $|x|$ is small, and we can truncate the series accordingly and add a bound for the omitted terms to the radius of the computed ball. When the series are truncated after the $k = K - 1$ term (for any $K < d + 1$), comparison with a geometric series shows that the error is bounded by the first omitted term times a simple factor.

PROPOSITION 7. *For $\sigma \in \{-1, +1\}$, the error when truncating the bottom sum in (24) or (25) (with prefactors removed) after the $k = K - 1$ term satisfies*

$$(27) \quad \left| \sum_{k=K}^d A_{\sigma}(d, k)(-x^2)^k \right| \leq \frac{A_{\sigma}(d, K)|x|^{2K}}{1-\alpha}, \quad \alpha = |x|^2 \frac{(d-K+1)(2d+2K+\sigma)}{K(2K+\sigma)}$$

provided that $\alpha < 1$.

For bounding $A_{\sigma}(d, K)$ in this expression, and for selecting the appropriate truncation point K , we use the binomial closed forms together with the remarks in [subsection 4.5](#).

The alternating series (24) and (25) may suffer from significant cancellation, which requires use of increased precision. We can estimate the magnitude by noting that no cancellation occurs if x is an imaginary number. Solving the majorizing recurrence $f_n = 2|z|f_{n-1} + f_{n-2}$ with $f_0 = 1$, $f_1 = |z|$ shows that

$$|P_n(z)| \leq |P_n(i|z)| \leq \left(|z| + \sqrt{1 + |z|^2} \right)^n.$$

Therefore, the possible cancellation assuming that $|P_n(x)| \approx 1$ is about

$$(28) \quad p_A = n \log_2(|x| + \sqrt{1 + |x|^2})$$

bits (which is at most $(\log_2(1 + \sqrt{2}))n \approx 1.27156n$), so using ball arithmetic with about $p + p_A$ bits of working precision for the series evaluation gives p -bit accuracy.

4.3. Expansion at one. Expanding at $x = 1$ yields

$$(29) \quad P_n(x) = \sum_{k=0}^n c_{n,k} u^k, \quad c_{n,k} = \binom{n}{k} \binom{n+k}{k}$$

where $u = (x - 1)/2$. The coefficients $c_{n,k}$ are hypergeometric with initial value $c_{n,0} = 1$ and term ratio

$$(30) \quad \frac{c_{n,k}}{c_{n,k-1}} = \frac{(n-k+1)(n+k)}{k^2}.$$

As in the previous section, we can truncate (29) and bound the error by comparison with a geometric series.

PROPOSITION 8. *The error when truncating (29) after the $k = K - 1$ term satisfies*

$$(31) \quad \left| \sum_{k=K}^n c_{n,k} u^k \right| \leq \frac{c_{n,K} |u|^K}{1 - \alpha}, \quad \alpha = |u| \frac{(n - K)(n + K + 1)}{(K + 1)^2}$$

provided that $\alpha < 1$.

For $u \geq 0$, (29) does not suffer from cancellation. For $u < 0$, we can estimate the amount of cancellation from the magnitude of $P_n(x')$ where $|u| = (x' - 1)/2$. For not too large $x' \geq 1$, a very good approximation is

$$P_n(x') \leq 2 \sum_{k=0}^{\infty} \frac{n^{2k}}{k!^2} |u|^k = 2 I_0(2n\sqrt{|u|}) \leq 2 e^{2n\sqrt{|u|}}.$$

We therefore need about $2n\sqrt{\max(0, -u)}/\ln(2)$ bits of increased precision.

We can compute P'_n from P_n and P_{n-1} , but since this involves a division by $1 - x^2$, it is useful to evaluate P'_n directly when x is close to 1. We have $P'_n(x) = \sum_{k=0}^{n-1} c'_{n,k} u^k$ where $c'_{n,k} = (k + 1)c_{n,k+1}/2$ satisfies

$$(32) \quad c'_{n,0} = \frac{n(n+1)}{2}, \quad \frac{c'_{n,k}}{c'_{n,k-1}} = \frac{(n-k)(n+k+1)}{k(k+1)}.$$

Since $c'_{n,k} \leq nc_{n,k+1}$, the analog

$$(33) \quad \left| \sum_{k=K}^n c'_{n,k} u^k \right| \leq n \binom{n}{K+1} \binom{n+K+1}{K+1} |u|^K \frac{1}{1-\alpha}$$

of Proposition 8 holds with u and α as above.

4.4. Hypergeometric series evaluation. We use rectangular splitting [23, 12] to evaluate hypergeometric series with rational parameters where the argument x is a high-precision number. This reduces evaluating a K -term series to $\mathcal{O}(K)$ cheap scalar operations (additions and multiplications or divisions by small integer coefficients) and about $2\sqrt{K}$ expensive nonscalar operations (general multiplications), whereas direct evaluation of the hypergeometric recurrence uses $\mathcal{O}(K)$ expensive operations.

Algorithm 2 presents our version of rectangular splitting for the present application. We implement the various series expansions by defining the polynomial evaluation functions $p(k), q(k)$ according to formulas (21), (26), (30), (32).

This algorithm is a generalization of the method for evaluating Taylor series of elementary functions given in [13], which combines rectangular splitting with partially unrolling the recurrence to reduce the number of scalar divisions (which in practice are more costly than scalar multiplications). The terms are computed in the reverse direction to allow using Horner's rule for the outer multiplications.

Our code uses ball arithmetic for x and s so that no error analysis is needed, and we use a bignum type for c (so no overflow can occur regardless of u). For low precision a faster implementation would be possible using fixed-point arithmetic with tight control of the word-level operations as was done for elementary functions in [13].

The algorithm contains two tuning parameters. The splitting parameter m controls the number m of multiplications for powers versus the number K/m of multiplications for Horner's rule. The choice $m \approx \sqrt{K}$ is optimal, but when evaluating two

Algorithm 2 Evaluation of hypergeometric series using rectangular splitting**Input:** An arbitrary x , recurrence data $p, q \in \mathbb{Z}[k]$, integer $K \geq 0$, offset $\Omega \in \{0, 1\}$.**Output:** $s = \sum_{k=\Omega}^{K-1} x^k \prod_{j=\Omega}^k p(j)/q(j)$

```

1:  $m \leftarrow \lfloor \sqrt{K} \rfloor$ ; precompute  $[1, x, x^2, \dots, x^m]$   $\triangleright$  Tuning parameter: any  $m \geq 1$  can be used
2:  $s \leftarrow 0$ ;  $k \leftarrow K - 1$ 
3: while  $k \geq \Omega$  do
4:    $u \leftarrow \min(4, k + 1 - \Omega)$   $\triangleright$  Tuning parameter: any  $1 \leq u \leq k + 1 - \Omega$  can be used
5:    $(a, b) \leftarrow (k - u + 1, k)$   $\triangleright$  Unrolled range
6:    $c \leftarrow \prod_{j=a}^b p(j)$   $\triangleright$  Small integer coefficient
7:   while  $k \geq a$  do
8:      $r \leftarrow k \bmod m$ 
9:     if  $k = b$  then
10:       $s \leftarrow c \cdot (s + x^r)$   $\triangleright$  Using precomputed power of  $x$ 
11:     else
12:       $s \leftarrow s + c \cdot x^r$   $\triangleright$  Using precomputed power of  $x$ 
13:     end if
14:     if  $r = 0$  and  $k \neq 0$  then
15:        $s \leftarrow s \cdot x^m$   $\triangleright$  Using precomputed power of  $x$ 
16:     end if
17:      $c \leftarrow (c/p(k))q(k)$   $\triangleright$  Exact small integer division
18:      $k \leftarrow k - 1$ 
19:   end while
20:    $s \leftarrow s/c$ 
21: end while
22: return  $s$ 

```

series for the same x (in our case, to compute both $P_n(x)$ and $P'_n(x)$), the table of powers can be reused, and then $m \approx \sqrt{2K}$ minimizes the total cost.

The unrolling parameter u controls the number of coefficients to collect on a single denominator, reducing the number of divisions to N/u . Ideally, u should be chosen so that $\prod_{j=a}^b p(j)$ and $\prod_{j=a}^b q(j)$ fit in 1 or 2 machine words. The example value $u = 4$ is a reasonable default, but as an optimization, one might vary u for each iteration of the main loop to ensure that c always fits in a specific number of words.

The redundant parameter Ω is a small convenience in the pseudocode. Setting $\Omega = 1$ and adding the constant term separately avoids having to make a special case to prevent division by zero when $q(0) = 0$.

Due to the scalar operations, rectangular splitting ultimately requires $\mathcal{O}(K)$ arithmetic operations with $\tilde{\mathcal{O}}(p)$ bit complexity each just like straightforward evaluation of the recurrence, so it is not a genuine asymptotic improvement, but it is an improvement in practice and can give more than a factor 100 speedup at very high precision. It is possible to genuinely reduce the complexity of evaluating a hypergeometric sequence to $\mathcal{O}(\sqrt{K} \log(K))$ arithmetic operations using a baby-step giant-step method that employs fast multipoint evaluation, but in practice rectangular splitting performs better until both K and p exceed 10^6 (see [12]).

Another technique for fast evaluation of hypergeometric series, binary splitting, would be useful when p is large and the argument x is a rational number with small numerator and denominator, but this case is not relevant for our application.

Binary splitting also forms the basis of the bit-burst method [7, Section 4], which permits evaluating any fixed hypergeometric series at any fixed point to absolute precision p in only $\tilde{\mathcal{O}}(p)$ bit operations. Computing $P_n(x)$ by this method would require

$\mathcal{O}(\log p)$ analytic continuation steps, each of which would entail two evaluations of general solutions of the Legendre differential equation (2), given by power series of radius of convergence $1 - |x|$ whose coefficients obey recurrences of order two. This is to be compared with a single series, given by a first-order recurrence and typically converging faster, for the expansions considered in subsection 4.1 to subsection 4.3. Thus, the bit-burst method is unlikely to be competitive in the range of precision we are interested in, especially when the asymptotic series can be used.

Nevertheless, it can be shown that the solutions with unit initial values at x_0 of (2) have Taylor coefficients c_k at x_0 bounded by $(n^2/(1 - |x_0|))^{\mathcal{O}(k)}$ uniformly in n and x_0 . As a consequence, the cost of computing any individual root of $P_n(x)$ by the bit-burst method and Newton iteration is $\tilde{\mathcal{O}}(p) \log(n)^{\mathcal{O}(1)}$. For computing all the roots, this approach matches the $\tilde{\mathcal{O}}(np)$ estimate of Theorem 1, while allowing for parallelization and requiring less memory. It may hence provide an alternative to multipoint evaluation worth investigating for precisions in the millions of bits.

4.5. Binomial coefficients. The prefactors of both the series expansion at $x = 0$ and the asymptotic series contain the central binomial coefficient $\binom{2n}{n}$. We need to compute this factor efficiently for any n and precision p . Since $\binom{2n}{n} \approx 4^n$, it is best to use an exact algorithm when $n < Cp$ for some small constant $C > 1/2$. We use the binomial function provided by GMP for $n < 6p + 200$ and otherwise use an asymptotic series for $\binom{2n}{n}$ with error bounds given in [5].

We also need to quickly estimate the magnitude of binomial coefficients for error bounds of series truncations. We have the binary entropy estimate

$$\log_2 \binom{n}{k} \leq nG(k/n), \quad G(x) = -x \log_2(x) - (1-x) \log_2(1-x)$$

and the equivalent form

$$(34) \quad \binom{n}{k} \leq \left(\frac{n}{n-k}\right)^{n-k} \left(\frac{n}{k}\right)^k = \frac{n^n}{k^k(n-k)^{n-k}}.$$

The function $G(x)$ can be evaluated cheaply with a precomputed lookup table. A coarse estimate is sufficient, since overestimating $\log_2 \binom{n}{k}$ by a few percent only adds a few percent to the running time.

5. Algorithm selection. We first use a set of cutoffs found experimentally to decide whether to use the basecase recurrence or one of the series expansions. The recurrence is mainly faster for some combinations of $p < 1000$, $n < 400$ when computing $(P_n(x), P'_n(x))$ simultaneously and for some combinations of $p < 500$, $n < 100$ when computing $P_n(x)$ alone (in all cases subject to some boundaries $\varepsilon < x < 1 - \varepsilon$); the actual optimal regions are complicated due to differences in overhead between fixed-point integer arithmetic and ball arithmetic for the respective algorithm implementations. For the actual cutoffs used, we refer to the source code.

To select between the series expansion at $x = 0$, the expansion at $x = 1$, and the asymptotic series, the following heuristic is used. For each algorithm A , we estimate the evaluation cost as $C_A = K_A(p + p_A)$ where K_A is the number of terms required by algorithm A ($K_A = \infty$ if A is the asymptotic series and it does not converge to the required accuracy), p is the precision goal, and p_A is the extra precision required by algorithm A due to internal cancellation. For the asymptotic series, we multiply the cost by an extra factor 2 as a penalty for using complex numbers. In the end, we select the algorithm with the lowest C_A .

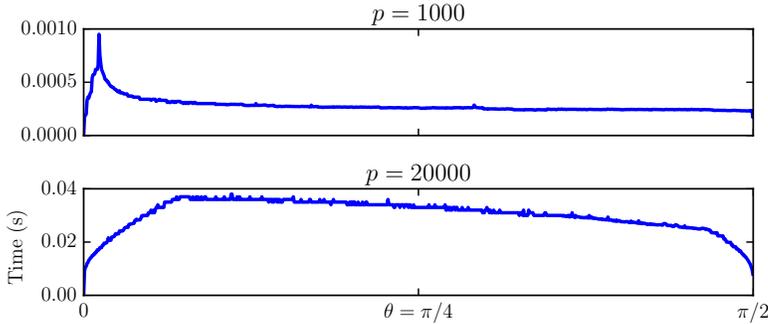


FIG. 1. Time to evaluate $(P_n(x), P'_n(x))$ as the argument $x = \cos(\theta)$ varies ($x = 1$ at $\theta = 0$ and $x = 0$ at $\theta = \pi/2$), here with $n = 10\,000$, for precision p somewhat smaller than n (top plot) and somewhat larger (bottom plot). The variable θ is used for the horizontal scale in this picture to follow the distribution of the roots (which are clustered near $x = 1$) linearly.

We select K_A and estimate p_A heuristically using machine precision floating-point computations, working with logarithmic magnitudes to avoid underflow and overflow. During the actual evaluation of the series expansions, K_A and p_A are then given; we compute rigorous upper bounds for the truncation error via (20), (22), (27), (31), (33), (34) using floating-point arithmetic with directed rounding, while additional rounding errors are tracked by the ball arithmetic.

The assumption that the running time is a bilinear function of K_A and $p + p_A$ is not completely realistic, but this cost estimate nonetheless captures the correct asymptotics when

$$x \rightarrow 0, \quad x \rightarrow 1, \quad n \rightarrow \infty, \quad p \rightarrow \infty$$

separately, and hopefully will not be too inaccurate in the transition regions. This is verified empirically.

Figure 1 illustrates how the time to evaluate $(P_n(x), P'_n(x))$ varies with x when the automatic algorithm selection is used. Here, we have timed the case $n = 10\,000$ for two different p . For large n and $p \ll n$ (top plot), a sharp peak appears at the transition between the series expansion at $x = 1$ and the asymptotic expansion which is used for most x . This peak tends to become taller but narrower for larger n . We could presumably get rid of the peak by implementing another algorithm specifically for the transition region, but the area under the peak is so small compared to the median baseline that computing all the roots would not be sped up much. For p somewhat larger than n (bottom plot), we observe a smooth transition between the series at $x = 1$ near the left of the picture and the series at $x = 0$ used over the most of the range.

6. Benchmarks. Except where otherwise noted, the timings were obtained on a 1.90 GHz Intel Core i5-4300U CPU using a single core.

6.1. Polynomial evaluation. Figure 2 compares the performance of different methods for evaluating $(P_n(x), P'_n(x))$ on a set of $n/2$ points distributed like the positive roots of $P_n(x)$ to simulate one stage of Newton iteration at p -bit precision. The time for the three-term recurrence is set to 1, i.e. we divide the other timings by this measurement. The following methods are timed:

- Our hybrid method (from here on called the “main algorithm”) with automatic selection between the three-term recurrence and different series expansions.

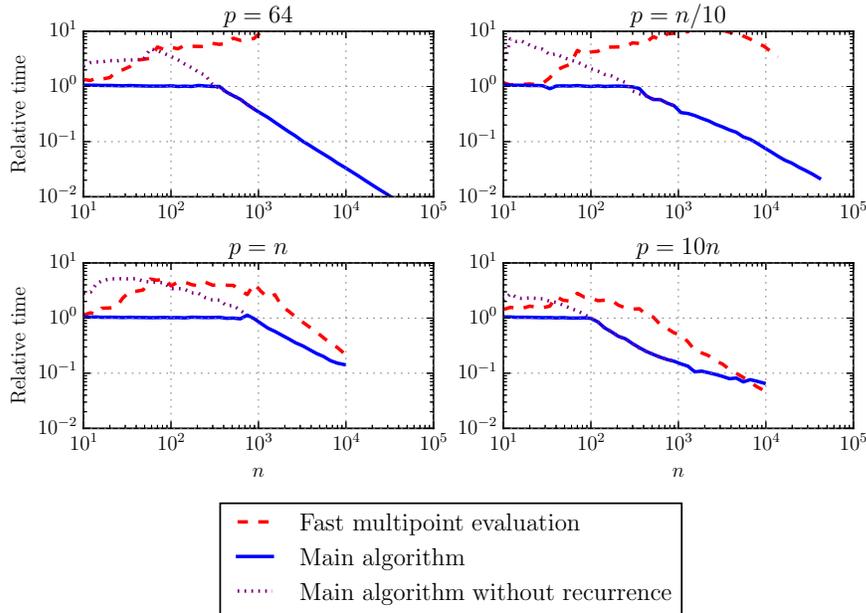


FIG. 2. Performance comparison of various methods to evaluate $(P_n(x), P'_n(x))$ to p -bit precision for a set of $n/2$ points $0 < x < 1$ distributed like the roots of P_n . The y axis (relative time) shows the time divided by the time using the three-term recurrence in fixed-point arithmetic.

- The main algorithm without the three-term recurrence as the basecase, i.e. using series expansions even for very small n .
- Fast multipoint evaluation of the expanded polynomials $P_n(x)$ and $P'_n(x)$. As in subsection 4.2, we expand $P_n(\sqrt{x})$ for even n and $P_n(\sqrt{x})/\sqrt{x}$ for odd n and evaluate at x^2 since this halves the amount of work. The polynomial coefficients are generated using the hypergeometric recurrence and the fast multipoint evaluation is done using `_arb_poly_evaluate_vec_fast_precomp` (where the “precomp” suffix indicates that the same product tree is used for both P_n and P'_n). The fast multipoint evaluation is done with $2.9n$ guard bits, which was found experimentally to be sufficient for full accuracy.

The crossover point between the three-term recurrence and series expansions usually occurs around $n \approx 10^2 - 10^3$ (it can be as low as $n \approx 10$ if p is much larger). For modest n , the three-term recurrence is much faster than the hypergeometric series (typically by a factor 3-4) due to working with negligible extra precision and thanks to the low overhead of fixed-point arithmetic. This low overhead is very useful for typical evaluation of Legendre polynomials and generation of quadrature nodes for one or a few machine words of precision. The crossover point could be lowered slightly if we used a similarly optimized fixed-point implementation for the hypergeometric series.

When p is fixed (top left in Figure 2), the main algorithm is a factor $\mathcal{O}(n)$ faster than the three-term recurrence since the asymptotic expansion converges to sufficient accuracy after $\mathcal{O}(1)$ terms for all sufficiently large n . With the constant precision $p = 64$, the main algorithm is 3.0 times faster for $n = 10^3$ and 30 times faster for $n = 10^4$. Conversely, fast multipoint evaluation with constant p is a factor $\mathcal{O}(n)$ slower than our algorithm due to the higher internal precision.

When $p \propto n$ (the three remaining plots in [Figure 2](#)), the main algorithm appears to show the same $\mathcal{O}(n)$ speedup over the three-term recurrence after the crossover point, at least initially. This speedup should level off asymptotically, but in practice this only occurs for n larger than 10^4 where we have already gained a factor 10 or more. The leveling off is visible in the bottom right figure ($p = 10n$).

Fast multipoint evaluation gives a true asymptotic $\mathcal{O}(n)$ speedup, but since it has much higher overhead, it only starts to give an improvement over the main algorithm from $n \approx 10^4$ and for p larger than n . When $p = n/10$, it appears that fast multipoint evaluation will only break even for n much larger than 10^5 . We conclude that fast multipoint evaluation would be worthwhile only for the last few Newton iterations when computing quadrature nodes for exceptionally high precision. Since independent evaluations are more convenient and easy to parallelize, the fast multipoint evaluation method currently seems to have limited practical value for this application.

6.2. Quadrature nodes. The function `arb_hypgeom_legendre_p_ui_root(x, w, n, k, p)` sets the output variable x to a ball containing the root of P_n with index k (we use the indexing $0 \leq k < n$, with $k = 0$ giving the root closest to 1), computed to p -bit precision. It also sets w to the corresponding quadrature weight. We use the formulas in [\[21\]](#) to compute an initial enclosure with roughly machine precision, followed by refinements with the interval Newton method at doubling precision steps for very high precision.

TABLE 1

Time in seconds to compute the set of degree- n Gauss-Legendre nodes and weights with p -bit precision. (For large n and p , the time was estimated by computing a subset of the nodes and weights.)

$n \setminus p$	64	256	1 024	3 333	33 333
20	0.000133	0.000229	0.000510	0.00121	0.0198
50	0.000450	0.000870	0.00212	0.00520	0.0710
100	0.00138	0.00310	0.00720	0.0163	0.191
200	0.00550	0.0111	0.0267	0.0550	0.589
500	0.0236	0.0610	0.164	0.325	2.61
1 000	0.0530	0.145	0.584	1.238	9.21
2 000	0.0860	0.298	1.12	4.20	32.6
5 000	0.191	0.665	2.67	14.3	181
10 000	0.350	1.26	4.93	26.6	674
100 000	3.60	12.2	41.3	212	13 637
1 000 000	58.0	146	411	1 850	103 960

[Table 1](#) shows absolute timings for computing degree- n Gauss-Legendre rules to p -bit precision by calling this function repeatedly with $0 \leq k < n/2$. [Table 2](#) shows the speedup ratio of our code compared to the `intnumgaussinit` function in Pari/GP which uses a generic polynomial root isolation strategy followed by Newton iteration for high precision refinement. The improvement is most dramatic for small p and large n where we benefit from using asymptotic expansions, but we also obtain a consistent speedup for large p .

For low precision and large n , our implementation is about three orders of magnitude slower than the machine precision code by Bogaert [\[3\]](#) which is reported to compute the nodes and weights for $n = 10^6$ in 0.02 seconds on four cores. This difference is reasonable since we use arbitrary-precision arithmetic, compute rigorous

TABLE 2

Speedup of our code compared to Pari/GP for generating degree- n Gauss-Legendre quadrature rules with p -bit precision.

$n \setminus p$	64	256	1 024	3 333	33 333
20	4.4	3.1	2.9	2.9	3.9
50	9.6	5.8	4.9	4.1	6.0
100	15	7.4	6.3	5.5	8.1
200	20	11	9.1	8.1	11
500	86	35	18	16	22
1 000	477	180	49	32	33
2 000	5 767	1 604	423	127	58

error bounds, and evaluate the Legendre polynomials explicitly whereas Bogaert uses a more sophisticated asymptotic development for both the nodes and the weights.

We also note that we can compute 53-bit floating-point values with provably correct rounding in about the same time as the 64-bit values, using Ziv’s strategy of increasing the precision. For a ball with relative radius just larger than 2^{-64} , there is less than a 1% probability that the correct 53-bit rounding cannot be determined, in which case that particular node can be recomputed with a few more bits.

Fousse [8] reports a few timings for smaller n and high precision obtained on a 2.40 GHz AMD Opteron 250 CPU. For example, $n = 80, p = 500$ takes 0.14 seconds (our implementation takes 0.029 seconds) and $n = 556, p = 5\,000$ takes 17 seconds (our implementation takes 0.53 seconds). Of course, these timings are not directly comparable since different CPUs were used.

The `mathinit` program included with version 2.2.19 of D. H. Bailey’s ARPREC library generates Gauss-Legendre quadrature nodes using Newton iteration together with the three-term recurrence for evaluating Legendre polynomials [2, 1]. With default parameters, this program computes the rules of degree $n = 3 \cdot 2^{i+1}$ for $1 \leq i \leq 10$ at 3 408 bits of precision, intended as a precomputation for performing degree-adaptive numerical integrations with up to 1 000 decimal digit accuracy. This takes about 1 300 seconds (our implementation takes 32 seconds). A breakdown for each degree level is shown in Table 3.

Table 3 also shows the approximation error and the evaluation time (using precomputed nodes and weights) for the degree- n approximations of three different integrals, illustrating the relative costs and realistic requirements for n . As motivation for the third integral, we might think of a segment of a Mellin-Barnes integral. The log, Airy and gamma function implementations in Arb are used.

The last few degree levels (with n roughly larger than the number of decimal digits) used by ARPREC tend to be dispensable for well-behaved integrands. A larger n is needed if the path of integration is close to a singularity or if the integrand is highly oscillatory. In such cases, bisecting the interval a few times to reduce the necessary n is often a better tradeoff. On the other hand, since the time to generate nodes with our code only grows linearly with n beyond $n \approx p$, increasing the degree further is viable, and potentially useful if the integrand is expensive to evaluate.

In the present work, we abstain from a more detailed discussion of adaptive integration strategies and the computation of error bounds for the integral itself. However, we mention that Arb contains an implementation of a version of the Petras algorithm [22] for rigorous integration. This code uses both adaptive path subdivision and Gauss-Legendre quadrature with an adaptive choice of n up to $n \approx 0.5p$

TABLE 3

Left columns: time in seconds to generate 1000-digit quadrature rules for the degrees n used by ARPREC. Right columns: for three different integrals, the error $|\int_{-1}^1 f(x)dx - \sum_{k=0}^{n-1} w_k f(x_k)|$ of the degree- n quadrature rule, and the time to evaluate this degree- n approximation of the integral at 1000-digit precision in Arb given precomputed nodes and weights (x_k, w_k) .

n	ARPREC	Our code	$\int_{-1}^1 \log(2+x)dx$		$\int_{-1}^1 \text{Ai}(10x)dx$		$\int_{-1}^1 \Gamma(1+ix)dx$	
			Error	Time	Error	Time	Error	Time
12	0.00520	0.000592	10^{-14}		10^{-1}		10^{-8}	
24	0.0189	0.00171	10^{-28}		10^{-9}		10^{-17}	
48	0.0629	0.00507	10^{-56}		10^{-34}		10^{-36}	
96	0.251	0.0163	10^{-111}		10^{-105}		10^{-73}	
192	0.974	0.0532	10^{-222}		10^{-284}	0.075	10^{-146}	
384	3.83	0.195	10^{-441}	0.023	10^{-721}	0.15	10^{-293}	1.3
768	15.2	0.763	10^{-881}	0.045	$< \varepsilon$	0.29	10^{-588}	2.5
1536	60.9	2.82	$< \varepsilon$	0.091			$< \varepsilon$	5.0
3072	241	9.55						
6144	1013	18.3						

TABLE 4

Left columns: step sizes h , number of evaluation points, and time to compute nodes for double exponential quadrature with Arb at 1000-digit precision. Right columns: error and evaluation time given precomputed nodes.

h	$2n+1$	Time	$\int_{-1}^1 \log(2+x)dx$		$\int_{-1}^1 \text{Ai}(10x)dx$		$\int_{-1}^1 \Gamma(1+ix)dx$	
			Error	Time	Error	Time	Error	Time
2^{-7}	1989	0.07	10^{-407}	0.12	10^{-423}	0.93	10^{-314}	6.3
2^{-8}	3977	0.14	10^{-814}	0.25	10^{-909}	1.75	10^{-630}	13.0
2^{-9}	7955	0.27	$< \varepsilon$	0.55	$< \varepsilon$	3.49	$< \varepsilon$	25.1

by default, with degree increments $n \approx 2^{k/2}$ and automatic caching of the nodes for fast repeated integrations. Node generation takes at most a few seconds for a first integration at 1000-digit precision and a few milliseconds for 100-digit precision.

7. Gauss-Legendre versus Clenshaw-Curtis and the double exponential method. The Clenshaw-Curtis and double exponential (tanh-sinh) quadrature schemes have received much attention as alternatives to Gauss-Legendre quadrature for numerical integration with very high precision [24, 1, 27]. Both schemes typically require a constant factor more evaluation points than Gauss-Legendre rules for equivalent accuracy, but the nodes and weights are easier to compute. Gauss-Legendre quadrature is therefore the best choice when the integrand is expensive to evaluate or when nodes can be precomputed for several integrations. It is of some interest to compare the relative costs empirically.

Here we assume an analytic integrand with singularities well isolated from the finite path of integration so that Gauss-Legendre quadrature is a good choice to begin with. As observed in [27], Clenshaw-Curtis often converges with identical rate to Gauss-Legendre for less well-behaved integrands, and the double exponential method is far superior to either Clenshaw-Curtis or Gauss-Legendre for analytic integrands with endpoint singularities.

Clenshaw-Curtis quadrature uses the Chebyshev nodes $\cos(\pi k/n), 0 \leq k \leq n$, and the corresponding weights can be expressed by a discrete cosine transform which

takes $\mathcal{O}(n \log n)$ arithmetic operations to compute by FFT. As a rule of thumb, $2n$ -point Clenshaw-Curtis quadrature gives the same accuracy as n -point Gauss-Legendre quadrature (for instance, the 384-point Clenshaw-Curtis rule gives errors of 10^{-229} , 10^{-294} and 10^{-154} for the three integrals in Table 3). As a point of comparison with Tables 1 and 3, Arb computes a length-2048 FFT with 1000-digit precision in 0.09 seconds and a length-32768 FFT with 10000-digit precision in 36 seconds. The precomputation for Clenshaw-Curtis quadrature is therefore roughly a factor 20 cheaper than for Gauss-Legendre quadrature with our algorithm, while subsequent integration with cached weights is twice as expensive for Clenshaw-Curtis.

Double exponential quadrature uses the change of variables $x = \tanh(\frac{1}{2}\pi \sinh t)$ to convert an integral on $(-1, 1)$ to the interval $(-\infty, +\infty)$ in such a way that the trapezoidal rule $\int_{-\infty}^{\infty} f(t) dt \approx h \sum_{k=-n}^n f(hk)$ converges exponentially fast. One generally chooses the discretization parameter as $h = 2^{-j}$ so that both the evaluation points and weights can be recycled for successive levels $j = 1, 2, 3 \dots$, and n is chosen so that the tail of the infinite series is smaller than 2^{-p} . The $2n+1$ nodes and weights can be computed with $n + \mathcal{O}(1)$ exponential function evaluations and $\mathcal{O}(n)$ arithmetic operations. Double exponential quadrature with Cn evaluation points typically achieves the same accuracy as n -point Gauss-Legendre quadrature, where C is slightly larger than for Clenshaw-Curtis, e.g. $C \approx 5$; see Table 4. The time to compute nodes and weights is comparable to Clenshaw-Curtis quadrature (around 0.2 seconds for 1000-digit precision and two minutes for 10000-digit precision).

In summary, for integration with precision in the neighborhood of 10^3 to 10^4 digits, computing n nodes with our algorithm is about an order of magnitude more expensive than performing n elementary function (e.g. exp or log) evaluations or computing an FFT of length n . This makes Gauss-Legendre quadrature competitive for computing more than m integrals (assuming that nodes and weights are cached), for a single integral requiring splitting into m subintervals, or for a single integral when the integrand costs more than m elementary function evaluations, where $m \approx 10^1$.

The picture becomes more complicated when accounting for the method used to estimate errors in an adaptive integration algorithm. One drawback of the Gauss-Legendre scheme is that the nodes are not nested, and therefore an adaptive strategy that repeatedly doubles the quadrature degree requires twice the number of function evaluations as the degree of the final level. However, this drawback disappears if the error is estimated by extrapolation or if an error bound is computed a priori as in the Petras algorithm [22].

8. Conclusion. In [1], it was claimed that “There is no known scheme for generating Gaussian abscissa–weight pairs that avoids [the] quadratic dependence on n . High-precision abscissas and weights, once computed, may be stored for future use. But for truly extreme-precision calculations – i.e., several thousand digits or more – the cost of computing them even once becomes prohibitive”.

In this quote, “quadratic dependence” refers to the number of arithmetic operations. We may remark that using asymptotic expansions in the evaluation of Legendre polynomials avoids the quadratic dependence on n for fixed precision p , and Theorem 1 avoids the implied cubic time dependence on n when $n \sim p$. In fact, Theorem 1 implies that Gauss-Legendre, Clenshaw-Curtis and double exponential quadrature have the same quasi-optimal asymptotic bit complexity, up to logarithmic factors.

Our experiments show that the algorithm in Theorem 1 hardly is worthwhile. However, our main algorithm does achieve a significant speedup for practical p and n which allows us to compute Gauss-Legendre quadrature rules for 1000-digit integra-

tion in 1-2 seconds and for 10 000-digit integration in 10-20 minutes on a single core. This is not prohibitively expensive compared to the repeated evaluation of typical integrands, especially if several integrations are needed. Parallelization is also trivial since all roots are computed independently.

A natural extension of this work would be to consider Gaussian quadrature rules for different weight functions. The techniques that work for Legendre polynomials should transfer to other classical orthogonal polynomials (Jacobi, Hermite, Laguerre, etc.) which likewise have hypergeometric expansions and satisfy three-term recurrence relations. The main obstacle might be to obtain large- n asymptotic expansions with suitable error bounds.

Acknowledgements. We thank Bruno Salvy for pointing out Petras' paper [21]. MM was supported in part by ANR grant ANR-14-CE25-0018-01 (FastRelax).

REFERENCES

- [1] D. H. BAILEY AND J. M. BORWEIN, *High-precision numerical integration: Progress and challenges*, Journal of Symbolic Computation, 46 (2011), pp. 741–754.
- [2] D. H. BAILEY, H. YOZO, X. S. LI, AND B. THOMPSON, *ARPREC: An arbitrary precision computation package*, Lawrence Berkeley National Laboratory, (2002).
- [3] I. BOGAERT, *Iteration-free computation of Gauss–Legendre quadrature nodes and weights*, SIAM Journal on Scientific Computing, 36 (2014), pp. A1008–A1026.
- [4] I. BOGAERT, B. MICHIELS, AND J. FOSTIER, *$O(1)$ computation of Legendre polynomials and Gauss–Legendre nodes and weights for parallel computing*, SIAM Journal on Scientific Computing, 34 (2012), pp. C83–C101.
- [5] R. P. BRENT, *Asymptotic approximation of central binomial coefficients with rigorous error bounds*, 2016. <http://arxiv.org/abs/1608.04834>.
- [6] Y. CHOW, L. GATTESCHI, AND R. WONG, *A Bernstein-type inequality for the Jacobi polynomial*, Proceedings of the American Mathematical Society, 121 (1994), pp. 703–709, <http://dx.doi.org/10.2307/2160265>.
- [7] D. V. CHUDNOVSKY AND G. V. CHUDNOVSKY, *Computer algebra in the service of mathematical physics and number theory*, in Computers in Mathematics, D. V. Chudnovsky and R. D. Jenks, eds., vol. 125 of Lecture Notes in Pure and Applied Mathematics, Dekker, 1990, pp. 109–232. Talks from the International Conference on Computers and Mathematics, Stanford University, 1986.
- [8] L. FOUSSE, *Accurate multiple-precision Gauss–Legendre quadrature*, in 18th IEEE Symposium on Computer Arithmetic, ARITH'07, IEEE, 2007, pp. 150–160.
- [9] G. H. GOLUB AND J. H. WELSCH, *Calculation of Gauss quadrature rules*, Mathematics of Computation, 23 (1969), pp. 221–230.
- [10] T. GRANLUND AND THE GMP DEVELOPMENT TEAM, *GNU MP: The GNU Multiple Precision Arithmetic Library*, 6.1.2 ed., 2017. <https://gmplib.org/>.
- [11] N. HALE AND A. TOWNSEND, *Fast and accurate computation of Gauss–Legendre and Gauss–Jacobi quadrature nodes and weights*, SIAM Journal on Scientific Computing, 35 (2013), pp. A652–A674.
- [12] F. JOHANSSON, *Evaluating parametric holonomic sequences using rectangular splitting*, in Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation, ISSAC '14, New York, NY, USA, 2014, ACM, pp. 256–263, <http://dx.doi.org/10.1145/2608628.2608629>.
- [13] F. JOHANSSON, *Efficient implementation of elementary functions in the medium-precision range*, in 22nd IEEE Symposium on Computer Arithmetic, ARITH22, 2015, pp. 83–89, <http://dx.doi.org/10.1109/ARITH.2015.16>.
- [14] F. JOHANSSON, *Computing hypergeometric functions rigorously*, 2016. <http://arxiv.org/abs/1606.06977>.
- [15] F. JOHANSSON, *Arb: efficient arbitrary-precision midpoint-radius interval arithmetic*, IEEE Transactions on Computers, 66 (2017), pp. 1281–1292, <http://dx.doi.org/10.1109/TC.2017.2690633>.
- [16] A. KOBEL AND M. SAGRALOFF, *Fast approximate polynomial multipoint evaluation and applications*, arXiv preprint arXiv:1304.8069, (2013).
- [17] M. A. KOWALSKI, A. G. WERSCHULZ, AND H. WOŹNIAKOWSKI, *Is Gauss quadrature optimal*

- for analytic functions?*, Numerische Mathematik, 47 (1985), pp. 89–98.
- [18] R. E. MOORE, *Methods and applications of interval analysis*, SIAM, 1979.
 - [19] F. W. J. OLVER, *Asymptotics and Special Functions*, A K Peters, Wellesley, MA, 1997.
 - [20] F. W. J. OLVER, D. W. LOZIER, R. F. BOISVERT, AND C. W. CLARK, *NIST Handbook of Mathematical Functions*, Cambridge University Press, New York, 2010.
 - [21] K. PETRAS, *On the computation of the Gauss–Legendre quadrature formula with a given precision*, Journal of Computational and Applied Mathematics, 112 (1999), pp. 253–267.
 - [22] K. PETRAS, *Self-validating integration and approximation of piecewise analytic functions*, Journal of Computational and Applied Mathematics, 145 (2002), pp. 345–359.
 - [23] D. M. SMITH, *Efficient multiple-precision evaluation of elementary functions*, Mathematics of Computation, 52 (1989), pp. 131–134, <http://myweb.lmu.edu/dmsmith/MComp1989.pdf>.
 - [24] H. TAKAHASI AND M. MORI, *Double exponential formulas for numerical integration*, Publications of the Research Institute for Mathematical Sciences, 9 (1974), pp. 721–741.
 - [25] THE PARI GROUP, *PARI/GP version 2.9.4*, Univ. Bordeaux, 2017. available from <http://pari.math.u-bordeaux.fr/>.
 - [26] A. TOWNSEND, *The race for high order Gauss–Legendre quadrature*, SIAM News, (2015), pp. 1–3.
 - [27] L. N. TREFETHEN, *Is Gauss quadrature better than Clenshaw–Curtis?*, SIAM review, 50 (2008), pp. 67–87.
 - [28] J. VAN DER HOEVEN, *Ball arithmetic*, tech. report, HAL, 2009. <http://hal.archives-ouvertes.fr/hal-00432152/fr/>.
 - [29] J. WIMP, *Computation with Recurrence Relations*, Pitman, Boston, 1984.