

The Borwein brothers, Pi and the AGM

Richard P. Brent

In fond memory of Jonathan M. Borwein 1951–2016

Abstract We consider some of Jonathan and Peter Borweins’ contributions to the high-precision computation of π and the elementary functions, with particular reference to their book *Pi and the AGM* (Wiley, 1987). Here “AGM” is the *arithmetic-geometric mean* of Gauss and Legendre. Because the AGM converges quadratically, it can be combined with fast multiplication algorithms to give fast algorithms for the n -bit computation of π , and more generally the elementary functions. These algorithms run in “almost linear” time $O(M(n)\log n)$, where $M(n)$ is the time for n -bit multiplication. We outline some of the results and algorithms given in *Pi and the AGM*, and present some related (but new) results. In particular, we improve the published error bounds for some quadratically and quartically convergent algorithms for π , such as the Gauss-Legendre algorithm. We show that an iteration of the Borwein-Borwein quartic algorithm for π is equivalent to two iterations of the Gauss-Legendre quadratic algorithm for π , in the sense that they produce exactly the same sequence of approximations to π if performed using exact arithmetic.

1 Introduction

Jonathan Borwein was fascinated by the constant π , and gave many stimulating talks on this topic. The slides for most of these talks may be found on the memorial website [11]. In my talk [21] at the Jonathan Borwein Commemorative Conference I discussed the reasons for this fascination. In a nutshell, it is that theorems about π are often just the tips of “mathematical icebergs” – much of interest lies hidden beneath the surface.

This paper considers some of Jonathan and Peter Borweins’ contributions to the high-precision computation of π and the elementary functions \log , \exp , \arctan , \sin , etc. The material is mainly drawn from their fascinating book *Pi and the AGM* [14].

Richard P. Brent

Mathematical Sciences Institute, Australian National University, Canberra, ACT 2601, Australia.
e-mail: JBCC@rpbrent.com

February 22, 2018

We make no attempt to review the whole book – a reader interested in the complete contents should consult one of the reviews [2, 3, 9, 47] or, better, read the book itself. We do not try to distinguish between the contributions of Jonathan and his brother Peter – so far as we know, they contributed equally to the book, although no doubt in different ways.

We take the opportunity to present some new results that are related to the material in *Pi and the AGM*. For example, the error after a finite number of iterations of some of the quadratically and quartically convergent algorithms for π can be expressed succinctly in terms of theta functions. Inspection of these expressions suggests that some algorithms, previously considered different, are actually equivalent, in the sense that they give exactly the same sequence of approximations to π if performed using exact arithmetic. For example, one of the Borweins’ quadratically convergent algorithms [14, Iteration 5.2 with $r = 4$] is equivalent to the Gauss-Legendre algorithm [17, 19, 41], and it follows that one step of the Borweins’ quartically convergent algorithm [14, Iteration 5.3] is equivalent to two steps of the Gauss-Legendre algorithm. These connections between superficially different algorithms do not seem to have been noticed before.

In §2 we give some necessary definitions, discuss the *arithmetic-geometric mean*, and consider its connection with elliptic integrals and Jacobi theta functions. We also mention the concept of *order of convergence* of an algorithm.

A brief history of quadratically convergent algorithms for π is given in §3.

In §4 we consider some quadratically and quartically convergent algorithms for π , including the Gauss-Legendre algorithm and several algorithms due to the Borweins. In §5 we show that some of the algorithms of §4, although superficially different, are actually equivalent when performed with exact arithmetic.

Chapter 5 of *Pi and the AGM* considers some striking *Ramanujan-Sato* formulæ for $1/\pi$ that give very fast (though linearly convergent) algorithms for computing π . The first such formulæ were given by Ramanujan [39]. Later authors include Takeshi Sato, the Borwein brothers, and the Chudnovsky brothers. See [6, 7, 15] for references. In §6 we briefly consider some Ramanujan-Sato formulæ and the corresponding algorithms for computing π .

One of the “icebergs” alluded to above is the fast computation of elementary functions to arbitrary precision. The constant $\pi = 4 \arctan(1)$ is of course just a special case (the tip of the iceberg). In §7 we outline how fast algorithms for computing elementary (and some other) functions can be based on the arithmetic-geometric mean iteration.

2 Preliminaries: means, elliptic integrals and theta functions

We define the *order of convergence* of a sequence. It will be sufficient to say that a sequence $(x_n)_{n \in \mathbb{N}}$ converges linearly to L (or with *order of convergence* 1) if

$$0 < \mu_0 = \liminf_{n \rightarrow \infty} \frac{|x_{n+1} - L|}{|x_n - L|} \leq \limsup_{n \rightarrow \infty} \frac{|x_{n+1} - L|}{|x_n - L|} = \mu_1 < 1.$$

If $\mu_0 = \mu_1$ then μ_0 is called the *rate of convergence*.

We say that a sequence $(x_n)_{n \in \mathbb{N}}$ *converges to L with order $p > 1$* if the sequence converges to L and there exists

$$p = \lim_{n \rightarrow \infty} \frac{\log |x_{n+1} - L|}{\log |x_n - L|} > 1.$$

Quadratic, cubic and quartic convergence are the cases $p = 2, 3, 4$ respectively. For example, if $x_n = 2^n \exp(-3^n)$, then $(x_n)_{n \in \mathbb{N}}$ converges cubically to zero, because $\log |x_{n+1}| / \log |x_n| = (-3^{n+1} + O(n)) / (-3^n + O(n)) \rightarrow 3$ as $n \rightarrow \infty$.

Roughly speaking, if a sequence converges linearly to L with rate μ , then the number of correct decimal digits in the approximation to L increases by about $\log_{10}(1/\mu)$ per term. For example, if

$$x_n = 2\sqrt{3} \sum_{j=0}^n \frac{(-1)^j}{(2j+1)3^j}, \quad (1)$$

then x_n converges linearly to π with about $\log_{10} 3 \approx 0.4771$ decimal digits per term.¹ If a sequence converges to L with order $p > 1$, then the number of correct digits is approximately multiplied by p for each additional term. For example, Newton's method for computing square roots²

$$x_{n+1} := \frac{1}{2} \left(x_n + \frac{S}{x_n} \right)$$

converges quadratically to $L := \sqrt{S}$, provided that x_0 and S are positive. In fact, it is easy to show that

$$x_{n+1} - L \approx \frac{1}{2L} (x_n - L)^2.$$

We now consider some well-known means. The *arithmetic mean* of $a, b \in \mathbb{R}$ is

$$\text{AM}(a, b) := \frac{a+b}{2},$$

and the *geometric mean* is

$$\text{GM}(a, b) := \sqrt{ab}. \quad (2)$$

Assuming that a and b are positive, we have the inequality

$$\text{GM}(a, b) \leq \text{AM}(a, b).$$

¹ The formula (1) is listed in Bailey's compendium [5], and is attributed to Madhava of Sangamagramma (c.1340–c.1425). It follows from the Taylor series for $\arctan(1/\sqrt{3})$.

² Attributed to Hero of Alexandria (c.10–70 A.D.), though also called the *Babylonian method*.

Initially we assume that a, b are positive real numbers. In §7 we permit a, b to be complex. To resolve the ambiguity in the square root in (2) we assume that $\Re(GM(a, b)) \geq 0$, and $\Im(GM(a, b)) \geq 0$ if $\Re(GM(a, b)) = 0$.

Given two positive reals a_0, b_0 , we can iterate the arithmetic and geometric means by defining, for $n \geq 0$,

$$\begin{aligned} a_{n+1} &= \text{AM}(a_n, b_n) \\ b_{n+1} &= \text{GM}(a_n, b_n). \end{aligned}$$

The sequences (a_n) and (b_n) converge quadratically to a common limit called the *arithmetic-geometric mean* (AGM) of a_0 and b_0 . We denote it by $\text{AGM}(a_0, b_0)$. Gauss [26] and Legendre [35] solved the problem of expressing $\text{AGM}(a, b)$ in terms of known functions. The answer may be written as

$$\frac{1}{\text{AGM}(a, b)} = \frac{2}{\pi} \int_0^{\pi/2} \frac{d\theta}{\sqrt{a^2 \cos^2 \theta + b^2 \sin^2 \theta}}. \quad (3)$$

The right-hand-side of (3) is the product of a constant (whose precise value will be significant later) and a *complete elliptic integral of the first kind*. As usual, the complete elliptic integral of the first kind is defined by

$$K(k) := \int_0^{\pi/2} \frac{d\theta}{\sqrt{1 - k^2 \sin^2 \theta}} = \int_0^1 \frac{dt}{\sqrt{(1-t^2)(1-k^2 t^2)}},$$

and the *complete elliptic integral of the second kind* by

$$E(k) := \int_0^{\pi/2} \sqrt{1 - k^2 \sin^2 \theta} d\theta = \int_0^1 \frac{\sqrt{1 - k^2 t^2}}{\sqrt{1 - t^2}} dt.$$

The variable k is called the *modulus*, and $k' := \sqrt{1 - k^2}$ is called the *complementary modulus*. It is customary to define

$$K'(k) := K(\sqrt{1 - k^2}) = K(k')$$

and

$$E'(k) := E(\sqrt{1 - k^2}) = E(k'),$$

so in the context of elliptic integrals a prime ($'$) does *not* denote differentiation. On the occasions when we need a derivative, we use operator notation

$$D_k K(k) := dK(k)/dk.$$

We remark that *Pi and the AGM* uses the “dot” notation $\dot{K}(k) := dK(k)/dk$, but this is potentially ambiguous and hard to see, so we prefer to avoid it.

The moduli k and k' can in general be complex, but unless otherwise noted we assume that they are real and in the interval $(0, 1)$.

In terms of the Gaussian hypergeometric function

$$F(a, b; c; z) := 1 + \frac{a \cdot b}{1! \cdot c} z + \frac{a(a+1) \cdot b(b+1)}{2! \cdot c(c+1)} z^2 + \dots,$$

we have

$$K(k) = \frac{\pi}{2} F\left(\frac{1}{2}, \frac{1}{2}; 1; k^2\right) \quad (4)$$

and

$$E(k) = \frac{\pi}{2} F\left(-\frac{1}{2}, \frac{1}{2}; 1; k^2\right). \quad (5)$$

From (4) and [1, 17.3.21], we also have³

$$K'(k) = \frac{2}{\pi} \log\left(\frac{4}{k}\right) K(k) - f(k), \quad (6)$$

where $f(k) = k^2/4 + O(k^4)$ is analytic in the disk $|k| < 1$.

Substituting $(a, b) \mapsto (1, k')$ in (3), and recalling that $k^2 + (k')^2 = 1$, we have

$$\text{AGM}(1, k') = \frac{\pi}{2K(k)}. \quad (7)$$

Thus, if we start from $a_0 = 1, b_0 = k' \in (0, 1)$ and apply the AGM iteration, $K(k)$ can be computed from

$$\lim_{n \rightarrow \infty} a_n = \frac{\pi}{2K(k)}. \quad (8)$$

$E(k)$ can be computed via the AGM at the same time as $K(k)$, using the well-known result [14, (b) on pg. 15]

$$\frac{E(k)}{K(k)} = 1 - \frac{k^2}{2} - \sum_{n=0}^{\infty} 2^n (a_n - a_{n+1})^2.$$

It follows from (4) and (6) that, for small k ,

$$K'(k) = (1 + O(k^2)) \log\left(\frac{4}{k}\right). \quad (9)$$

This will be relevant in §7. A bound on the $O(k^2)$ term is given in [14, Thm. 7.2].

The Gauss-Legendre algorithm depends on *Legendre's relation*: for $0 < k < 1$,

$$E(k)K'(k) + E'(k)K(k) - K(k)K'(k) = \frac{\pi}{2}.$$

For a proof, see *Pi and the AGM*, Sec. 1.6.

A computationally important special case, obtained by taking $k = k' = 1/\sqrt{2}$, is

$$\left(2E(1/\sqrt{2}) - K(1/\sqrt{2})\right) K(1/\sqrt{2}) = \frac{\pi}{2}. \quad (10)$$

³ Here and elsewhere, \log denotes the natural logarithm.

It can be shown [14, Thm. 1.7] that the two factors in (10) are

$$K(1/\sqrt{2}) = \frac{\Gamma^2(\frac{1}{4})}{4\pi^{1/2}} \quad \text{and} \quad 2E(1/\sqrt{2}) - K(1/\sqrt{2}) = \frac{\Gamma^2(\frac{3}{4})}{\pi^{1/2}}.$$

To estimate the order of convergence and to obtain error bounds, we consider the parameterisation of the AGM in terms of *Jacobi theta functions*. We need the basic theta functions of one variable, defined for $|q| < 1$ by

$$\theta_2(q) := \sum_{n \in \mathbb{Z}} q^{(n+1/2)^2}, \quad \theta_3(q) := \sum_{n \in \mathbb{Z}} q^{n^2}, \quad \theta_4(q) := \sum_{n \in \mathbb{Z}} (-1)^n q^{n^2}.$$

The theta functions satisfy many identities [46, §21.3]. In particular, we use the following addition formulæ, due to Jacobi [31]. They are proved in [14, §2.1].

$$\theta_3^2(q) = \theta_2^2(q^2) + \theta_3^2(q^2), \quad (11)$$

$$\theta_3^4(q) = \theta_2^4(q) + \theta_4^4(q). \quad (12)$$

It is not difficult to show that

$$\frac{\theta_3^2(q) + \theta_4^2(q)}{2} = \theta_3^2(q^2) \quad \text{and} \quad \sqrt{\theta_3^2(q)\theta_4^2(q)} = \theta_4^2(q^2).$$

Thus, the AGM variables (a_n, b_n) can, if scaled suitably, be parameterised by $(\theta_3^2(q^{2^n}), \theta_4^2(q^{2^n}))$. More precisely, if $1 = a_0 > b_0 = \theta_4^2(q)/\theta_3^2(q) > 0$, where $q \in (0, 1)$, then the variables a_n, b_n appearing in the AGM iteration satisfy

$$a_n = \frac{\theta_3^2(q^{2^n})}{\theta_3^2(q)}, \quad b_n = \frac{\theta_4^2(q^{2^n})}{\theta_3^2(q)}. \quad (13)$$

It is useful to define auxiliary variables $c_{n+1} := a_n - a_{n+1} = (a_n - b_n)/2$. Using the quotient for a_n and the addition formula (11), we see that

$$c_n = \frac{\theta_2^2(q^{2^n})}{\theta_3^2(q)} \quad (14)$$

holds for $n \geq 1$. We could use (14) to define c_0 , but this will not be necessary.⁴

We can write q (which is called the *nome*) explicitly, in fact

$$q = \exp(-\pi K'(k)/K(k)). \quad (15)$$

This is due to Gauss/Jacobi; for a proof see [14, Thm. 2.3]. In the special case $k = k' = 1/\sqrt{2}$, we have $K' = K$ and $q = e^{-\pi} = 0.0432139\dots$

⁴ Salamin [41] defines c_n using the relation $c_n^2 = a_n^2 - b_n^2$. This has the advantage that c_0 is defined naturally, and for $n > 0$ it is equivalent to our definition. However, it is computationally more expensive to compute $(a_n^2 - b_n^2)^{1/2}$ than $a_n - a_{n+1}$.

Because the AGM iteration converges quadratically, it offers the prospect of quadratically convergent algorithms for approximating π and, more generally, all the elementary functions. This is the topic of §4 and §7 below. First we make some comments on the history of quadratically convergent algorithms for π .

3 Historical remarks

An algorithm for computing $\log(4/k)$, using (7), (9) and the AGM, assuming that we know π to sufficient accuracy, was given by Salamin [8, pg. 71] in 1972. On the same page Salamin gives an algorithm for computing π , taking $k = 4/e^n$ in (9). With his choice $\pi \approx 2n \text{AGM}(1, k)$. However, this assumes that we know e , so it is not a “standalone” algorithm for π via the AGM. Similarly, if we take $k = 4/2^n$ in (9), we obtain an algorithm for computing $\pi \log 2$ (and hence π , if we know $\log 2$).

In 1975, Salamin [41] and (independently) the present author [17, 19] discovered a quadratically convergent algorithm for computing π via the AGM *without* needing to know e or $\log 2$ to high precision. It is known as the “Gauss-Legendre” algorithm (after the discoverers of the key identities [25, 35]) or the “Brent-Salamin” algorithm (after the 20th century discoverers [20]), and is about twice as fast as the earlier algorithms which assume a knowledge of e or $\log 2$. We abbreviate the name to *Algorithm GL*. Bailey and Borwein, in *Pi: The Next Generation* [6, Synopsis of paper 1], say “This remarkable co-discovery arguably launched the modern computer era of the computation of π ”.⁵

In 1984, Jon and Peter Borwein [12] (see also [14, Alg. 2.1]) discovered another quadratically convergent algorithm for computing π , with convergence about as fast as Algorithm GL. We call this the (first) *Borwein-Borwein algorithm*, abbreviated *Algorithm BB1*. Yet another quadratically convergent algorithm, which we call the (second) *Borwein-Borwein algorithm* and abbreviate as *Algorithm BB2*, dates from 1986 – see [13] and [14, Iteration 5.1]. Although Algorithm BB2 appears different from Algorithm GL, we show in §5 that the two algorithms are in fact equivalent, in the sense of producing the same sequence of approximations to π . This surprising fact does not seem to have been noticed before.

4 Some superlinearly convergent algorithms for π

In this section we describe the Gauss-Legendre algorithm (GL) and two quadratically convergent algorithms (BB1 and BB2) due to Jon and Peter Borwein. We also describe a 4-th order algorithm (BB4) due to the Borweins.

⁵ In [10, §10], Jon Borwein says “It [Algorithm GL] is based on the arithmetic-geometric mean iteration (AGM) and some other ideas due to Gauss and Legendre around 1800, although neither Gauss, nor many after him, ever directly saw the connection to effectively computing π ”.

Using Legendre's relation and the formulæ that we have given for E and K in terms of the AGM iteration, it is not difficult to derive Algorithm GL. We present it in pseudo-code using the same style as the algorithms in [22].

Algorithm GL

Input: The number of iterations n_{max} .

Output: A sequence of n_{max} intervals containing π .

$a_0 := 1; b_0 := 1/\sqrt{2}; s_0 := \frac{1}{4}$.

for n from 0 to $n_{max} - 1$ **do**

$$a_{n+1} := (a_n + b_n)/2;$$

$$c_{n+1} := a_n - a_{n+1};$$

output $(a_{n+1}^2/s_n, a_n^2/s_n)$.

if $n < n_{max} - 1$ **then**

$$b_{n+1} := \sqrt{a_n b_n};$$

$$s_{n+1} := s_n - 2^n c_{n+1}^2.$$

Remarks

1. Subscripts on variables such as a_n, b_n are given for expository purposes. In an efficient implementation only a constant number of real variables are needed, because a_{n+1} can overwrite a_n (after saving a_n in a temporary variable for use in the computation of b_{n+1}), and similarly for b_n, c_n and s_n .
2. The purpose of the final "if ... then" is simply to avoid unnecessary computations after the final output. Similar comments apply to the other algorithms given below.
3. Salamin [41] notes the identity $4a_{n+1}c_{n+1} = c_n^2$ which can be used to compute c_{n+1} without the numerical cancellation that occurs when using the definition $c_{n+1} = a_n - a_{n+1}$. However, this refinement costs time and is unnecessary, because the terms $2^n c_{n+1}^2$ diminish rapidly and make only a minor contribution to the overall error caused by using finite-precision real arithmetic. To obtain an accurate result it is sufficient to use $O(\log n_{max})$ guard digits.

Neglecting the effect of rounding errors, Algorithm GL gives a sequence of lower and upper bounds on π :

$$\frac{a_{n+1}^2}{s_n} < \pi < \frac{a_n^2}{s_n},$$

and both bounds converge quadratically to π . The lower bound is more accurate, so the algorithm is often stated with just the lower bound a_{n+1}^2/s_n (we call this variant *Algorithm GLI*). Table 1 shows the approximations to π given by the first few iterations. Correct digits are shown in bold. The quadratic convergence is evident.

Recall that in Algorithm GL we have $a_0 = 1, b_0 = 1/\sqrt{2}, s_0 = \frac{1}{4}$ and, for $n \geq 0$,

$$a_{n+1} = \frac{a_n + b_n}{2}, \quad b_{n+1} = \sqrt{a_n b_n}, \quad c_{n+1} = a_n - a_{n+1}, \quad s_{n+1} = s_n - 2^n c_{n+1}^2.$$

n	lower bound a_{n+1}^2/s_n	upper bound a_n^2/s_n
0	2.914213562373095048801689	$< \pi < 4.000000000000000000000000$
1	3.140579250522168248311331	$< \pi < \mathbf{3.187672642712108627201930}$
2	3.141592646213542282149344	$< \pi < \mathbf{3.141680293297653293918070}$
3	3.141592653589793238279513	$< \pi < \mathbf{3.141592653895446496002915}$
4	3.141592653589793238462643	$< \pi < \mathbf{3.141592653589793238466361}$

Table 1: Convergence of Algorithm GL

Take $q = e^{-\pi}$, and write

$$a_\infty := \lim_{n \rightarrow \infty} a_n = \theta_3^{-2}(q) = 2\pi^{3/2}/\Gamma^2(\frac{1}{4}) \approx 0.8472, \quad (16)$$

$$s_\infty := \lim_{n \rightarrow \infty} s_n = \theta_3^{-4}(q)/\pi = 4\pi^2/\Gamma^4(\frac{1}{4}) \approx 0.2285. \quad (17)$$

Since $c_n = \theta_2^2(q^{2^n})/\theta_3^2(q)$, we have

$$s_n - s_\infty = \theta_3^{-4}(q) \sum_{m=n}^{\infty} 2^m \theta_2^4(q^{2^{m+1}}). \quad (18)$$

Write $a_n/a_\infty = 1 + \delta_n$ and $s_n/s_\infty = 1 + \varepsilon_n$. Then

$$\delta_n = \theta_3^2(q^{2^n}) - 1 \sim 4q^{2^n} \text{ as } n \rightarrow \infty,$$

and (17) – (18) give

$$\varepsilon_n = \pi \sum_{m=n}^{\infty} 2^m \theta_2^4(q^{2^{m+1}}) \sim 2^{n+4} \pi q^{2^{n+1}}.$$

Writing

$$\frac{a_n^2/a_\infty^2}{s_n/s_\infty} = \frac{a_n^2}{\pi s_n} = \frac{(1 + \delta_n)^2}{1 + \varepsilon_n},$$

it is straightforward to obtain an upper bound on π :

$$0 < a_n^2/s_n - \pi < U(n) := 8\pi q^{2^n}. \quad (19)$$

Convergence is quadratic: if $e_n := a_n^2/s_n - \pi$, then

$$\lim_{n \rightarrow \infty} e_{n+1}/e_n^2 = \frac{1}{8\pi}.$$

Replacing a_n by a_{n+1} and δ_n by δ_{n+1} , we obtain a lower bound on π :

$$0 < \pi - \frac{a_{n+1}^2}{s_n} < L(n) := (2^{n+4} \pi^2 - 8\pi) q^{2^{n+1}}. \quad (20)$$

Pi and the AGM [(2.5.7) on page 48] gives a slightly weaker lower bound which, via (16), may be written as

$$\pi - \frac{a_{n+1}^2}{s_n} \leq \frac{2^{n+4}\pi^2 q^{2^{n+1}}}{a_\infty^2}. \quad (21)$$

Since $a_\infty^2 < 1$, the bound (21) is weaker than the bound (20). In (20), the factor $(2^{n+4}\pi^2 - 8\pi)$ is the best possible, since an expansion of a_{n+1}^2/s_n in powers of q gives $\pi - a_{n+1}^2/s_n = (2^{n+4}\pi^2 - 8\pi)q^{2^{n+1}} - O(2^n q^{2^{n+2}})$, with the minus sign before the “ O ” term informally indicating the sign of the remainder.

n	$a_n^2/s_n - \pi$	$\pi - a_{n+1}^2/s_n$	$\frac{a_n^2/s_n - \pi}{U(n)}$	$\frac{\pi - a_{n+1}^2/s_n}{L(n)}$
0	8.58e-1	2.27e-1	0.790369040	0.916996189
1	4.61e-2	1.01e-3	0.981804947	0.999656206
2	8.76e-5	7.38e-9	0.999922813	0.999999998
3	3.06e-10	1.83e-19	0.999999999	1.000000000
4	3.72e-21	5.47e-41	1.000000000	1.000000000
5	5.50e-43	2.41e-84	1.000000000	1.000000000
6	1.20e-86	2.31e-171	1.000000000	1.000000000
7	5.76e-174	1.06e-345	1.000000000	1.000000000
8	1.32e-348	1.11e-694	1.000000000	1.000000000

Table 2: Numerical values of upper and lower bounds for Algorithm GL

In Table 2, $U(n) := 8\pi \exp(-2^n \pi)$ and $L(n) := (2^{n+4}\pi^2 - 8\pi) \exp(-2^{n+1} \pi)$ are the bounds given in (19)–(20). It can be seen that the bounds are very accurate for $n > 1$, as expected from our analysis.

Recall that Algorithm GL gives approximations a_n^2/s_n and a_{n+1}^2/s_n to $\pi = a_\infty^2/s_\infty$. Using the expressions for a_n and s_n in terms of theta functions, we see that

$$\pi = \frac{a_n^2 \theta_3^{-4}(q^{2^n})}{s_n - \theta_3^{-4}(q) \sum_{m=n}^{\infty} 2^m \theta_2^4(q^{2^{m+1}})}, \quad (22)$$

[or similarly with the numerator replaced by $a_{n+1}^2 \theta_3^{-4}(q^{2^{n+1}})$]. The expression (22) for π is essentially of the form

$$\pi = \frac{a_n^2 - O(q^{2^n})}{s_n - O(2^n q^{2^{n+1}})} \left[\text{or } \frac{a_{n+1}^2 - O(q^{2^{n+1}})}{s_n - O(2^n q^{2^{n+1}})} \right].$$

This shows precisely how Algorithm GL approximates π and why it provides upper [or lower] bounds.

In *Pi and the AGM*, Jon and Peter Borwein present a quadratically convergent algorithm for π , based on the AGM, but different from Algorithm GL. It is Algorithm 2.1 in Chapter 2, and was first published in [12]. We call it *Algorithm BB1*.

Instead of using Legendre's relation, Algorithm BB1 uses the identity

$$K(k) D_k K(k) \Big|_{k=1/\sqrt{2}} = \frac{\pi}{\sqrt{2}},$$

where D_k denotes differentiation with respect to k .

Using the connection between $K(k')$ and the AGM, the Borweins [14, (2.4.7)] prove that

$$\pi = 2^{3/2} \frac{(\text{AGM}(1, k'))^3}{D_k \text{AGM}(1, k')} \Big|_{k=1/\sqrt{2}}.$$

An algorithm for approximating the derivative in this formula can be obtained by differentiating the AGM iteration symbolically. Details are given in [14].

We now present Algorithm BB1. Note that the algorithm given in [14] defines the upper bound $\bar{\pi}_n := \bar{\pi}_{n-1}(x_n + 1)/(y_n + 1)$ and omits the lower bound $\underline{\pi}_n$, but $\underline{\pi}_n$ can be obtained from [14, ex. 2.5.11]. We present a version that computes upper ($\bar{\pi}_n$) and lower ($\underline{\pi}_n$) bounds for comparison with Algorithm GL.

Algorithm BB1

Input: The number of iterations n_{max} .

Output: A sequence of n_{max} intervals containing π .

$x_0 := \sqrt{2}$;

output ($\underline{\pi}_0 := x_0, \bar{\pi}_0 := x_0 + 2$).

$x_1 := x_0^{1/2}$; $x_1 := \frac{1}{2}(x_0^{1/2} + x_0^{-1/2})$;

for n from 1 to $n_{max} - 1$ **do**

$$\underline{\pi}_n := \frac{2\bar{\pi}_{n-1}}{y_n + 1}; \bar{\pi}_n := \underline{\pi}_n \left(\frac{x_n + 1}{2} \right);$$

output ($\underline{\pi}_n, \bar{\pi}_n$);

if $n < n_{max} - 1$ **then**

$$x_{n+1} := \frac{1}{2}(x_n^{1/2} + x_n^{-1/2}); y_{n+1} := \frac{y_n x_n^{1/2} + x_n^{-1/2}}{y_n + 1}.$$

It may be shown that $\bar{\pi}_n$ decreases monotonically to the limit π , and $\underline{\pi}_n$ increases monotonically to π . Moreover, $\bar{\pi}_n - \underline{\pi}_n$ decreases quadratically to zero. This is illustrated in Table 3.

It is not immediately obvious that Algorithm BB1 depends on the AGM. However, the AGM is present in *Legendre form*: if $a_0 := 1, b_0 := k' = 1/\sqrt{2}$, and we perform n steps of the AGM iteration to define a_n, b_n , then $x_n = a_n/b_n$ and, for $n \geq 1, y_n = D_k b_n / D_k a_n$.

Comparing Tables 1 and 3, we see that Algorithm BB1 gives better upper bounds, but worse lower bounds, than Algorithm GL, for the same value of n (i.e. same number of square roots).

As for Algorithm GL, we can express the error after n iterations of Algorithm BB1 using theta functions, and deduce the asymptotic behaviour of the error.

Consider the AGM iteration with $a_0 = 1, b_0 = k' = (1 - k^2)^{1/2}$. Then a_n and b_n are functions of k . In *Pi and the AGM* it is shown that, for $n \geq 1$,

n	$\underline{\pi}_n$	$\bar{\pi}_n$
0	1.414213562373095048801689	$< \pi < 3.414213562373095048801689$
1	3.119132528827772757303373	$< \pi < \mathbf{3.142606753941622600790720}$
2	3.141548837729436193482357	$< \pi < \mathbf{3.141592660966044230497752}$
3	3.141592653436966609787790	$< \pi < \mathbf{3.141592653589793238645774}$
4	3.141592653589793238460785	$< \pi < \mathbf{3.141592653589793238462643}$

Table 3: Convergence of Algorithm BB1

$$\bar{\pi}_{n-1} = \left(2^{3/2} b_n^2 a_n / D_k a_n \right) \Big|_{k=1/\sqrt{2}}. \quad (23)$$

Now a_n and b_n are given by (13) with $q = e^{-\pi}$. We differentiate a_n with respect to k , where $k = (1 - b_0^2)^{1/2} = \theta_2^2(q) / \theta_3^2(q)$. This gives

$$D_k a_n = D_q \left(\frac{\theta_3^2(q^{2^n})}{\theta_3^2(q)} \right) / D_q \left(\frac{\theta_2^2(q)}{\theta_3^2(q)} \right) \Big|_{q=e^{-\pi}}. \quad (24)$$

We remark that (24) gives $D_k a_0 = 0$, as expected since a_0 is independent of k .

Thanks to the analyticity of the theta functions in $|q| < 1$, there is no difficulty in showing that⁶

$$\lim_{n \rightarrow \infty} D_k a_n = D_k \lim_{n \rightarrow \infty} a_n.$$

We denote the common value by $D_k a_\infty$. Taking the limit in (23), we obtain (as also follows from [14, (2.4.7)]):

$$D_k a_\infty = \frac{2^{3/2} a_\infty^3}{\pi} = 0.547486\dots \quad (25)$$

Now $a_n - a_\infty = \sum_{m=n+1}^{\infty} c_m$, and differentiating both sides with respect to k gives

$$D_k a_n - D_k a_\infty = \sum_{m=n+1}^{\infty} D_q \left(\frac{\theta_2^2(q^{2^m})}{\theta_3^2(q)} \right) / D_q \left(\frac{\theta_2^2(q)}{\theta_3^2(q)} \right) \Big|_{q=e^{-\pi}}. \quad (26)$$

We remark that (26) is analogous to (18), which we used in the analysis of Algorithm GL. Using (23) – (26), we obtain an upper bound on π (for $n \geq 1$, $q = e^{-\pi}$)

$$0 < \bar{\pi}_n - \pi < 2^{n+4} \pi^2 q^{2^{n+1}}. \quad (27)$$

A slightly weaker bound than (27) is proved in [14, §2.5].

⁶ Similarly, where we exchange the order of taking derivatives and limits elsewhere in this section, it is easy to justify.

Similarly, we can obtain a lower bound on π :

$$0 < \pi - \underline{\pi}_n < 4\pi q^{2^n}. \quad (28)$$

We omit detailed proofs of (27) and (28); they involve straightforward but tedious expansions of power series in q . Experimental evidence is provided in Table 4.

n	$\bar{\pi}_n - \pi$	$\frac{\bar{\pi}_n - \pi}{2^{n+4}\pi^2 q^{2^{n+1}}}$	$\pi - \underline{\pi}_n$	$\frac{\pi - \underline{\pi}_n}{4\pi q^{2^n}}$
1	1.01e-3	0.9896487063	2.25e-2	0.9570949132
2	7.38e-9	0.9948470082	4.38e-5	0.9998316841
3	1.83e-19	0.9974691480	1.53e-10	0.9999999988
4	5.47e-41	0.9987456847	1.86e-21	1.0000000000
5	2.41e-84	0.9993755837	2.75e-43	1.0000000000
6	2.31e-171	0.9996884727	6.01e-87	1.0000000000
7	1.06e-345	0.9998444059	2.88e-174	1.0000000000
8	1.11e-694	0.9999222453	6.59e-349	1.0000000000

Table 4: Numerical values of upper and lower bounds for Algorithm BB1

Table 4 gives numerical values of the approximation errors $\bar{\pi}_n - \pi$ and $\pi - \underline{\pi}_n$, and the ratio of these values to the bounds (27) and (28) respectively. It can be seen that the bounds are very accurate (as expected from the expressions for the errors in terms of theta functions and the rapid convergence of the series for the theta functions). The upper bound overestimates the error by a factor of $1 + O(2^{-n})$. A computation shows that we can not replace the bound by the function $L(n)$ defined in (20), although a similar bound appears to be valid if the constant 8π in (20) is replaced by a slightly smaller constant, e.g. 7π .

The bounds (27)–(28) can be compared with the lower bound $(2^{n+4}\pi^2 - 8\pi)q^{2^{n+1}}$ and upper bound $8\pi q^{2^n}$ for Algorithm GL. The upper bound is better for Algorithm BB1, but the lower bound is better for Algorithm GL. This confirms the observation above regarding the comparison of Tables 1 and 3.

Since it will be needed in §5, we state another quadratic algorithm, *Algorithm BB2*, different from Algorithm BB1 but also due to Jon and Peter Borwein (iteration 5.2 on page 170 of [14] with the parameter $r = 4$).

Algorithm BB2**Input:** The number of iterations n_{max} .**Output:** A sequence of n_{max} approximations to π . $\alpha_0 := 6 - 4\sqrt{2}; k_0 := 3 - 2\sqrt{2};$ **for** n from 0 to $n_{max} - 1$ **do****output** $\hat{\pi}_n := 1/\alpha_n;$ **if** $n < n_{max} - 1$ **then**

$$k'_n := \sqrt{1 - k_n^2}; k_{n+1} := \frac{1 - k'_n}{1 + k'_n};$$

$$\alpha_{n+1} := (1 + k_{n+1})^2 \alpha_n - 2^{n+2} k_{n+1}.$$

In Algorithm BB2, we have $\hat{\pi}_n \rightarrow \pi$ quadratically [14, pg. 170]. We remark that it would be clearer to increase (by one) the subscripts on the variables in Algorithm BB2, so as to correspond to the usage in Algorithm GL, which implicitly has $k'_0 = b_0/a_0 = 1/\sqrt{2}$ and $k_1 = (1 - k'_0)/(1 + k'_0) = 3 - 2\sqrt{2}$, but we have kept the notation used in [14].

The Borwein brothers did not stop at quadratic (second-order) algorithms for π . In Chapter 5 of *Pi and the AGM* they gave algorithms of orders 3, 4, 5 and 7. Of course, these algorithms are not necessarily faster than the quadratic algorithms, because we must take into account the amount of work per iteration. For a fair comparison, we can use Ostrowski's *efficiency index* [38, §3.11], defined as $\log(p)/W$, where $p > 1$ is the order of convergence and W is the work per iteration. A justification of this measure of efficiency is given in [16]. Consider a simple example – if we combine three iterations of Algorithm BB2 into one iteration of a new algorithm, then we obtain an algorithm of order 8, but with three times as much work per iteration. The efficiency index is the same in both cases, as it should be.

We refer to [14, Chapter 5] for the Borweins' cubic, quintic and higher-order algorithms, and consider only their quartic algorithm, which we call *Algorithm BB4*. It is a specialisation to the case $r = 4$ of the slightly more general algorithm given in [14, iteration 5.3, pg. 170]. The same special case is given in [15, Algorithm 1] and has been used in extensive calculations of π , see for example [4, 32]. We have changed notation slightly ($a_n \mapsto z_n$) to avoid conflict with the notation used in Algorithm GL.

Algorithm BB4**Input:** The number of iterations n_{max} .**Output:** A sequence of n_{max} approximations to π . $y_0 := \sqrt{2} - 1$; $z_0 := 2y_0^2$;**for** n from 0 to $n_{max} - 1$ **do** **output** $\pi_n := 1/z_n$; **if** $n < n_{max} - 1$ **then**

$$y_{n+1} := \frac{1 - (1 - y_n^4)^{1/4}}{1 + (1 - y_n^4)^{1/4}};$$

$$z_{n+1} := z_n(1 + y_{n+1})^4 - 2^{2n+3}y_{n+1}(1 + y_{n+1} + y_{n+1}^2).$$

In Algorithm BB4, π_n converges quartically to π . A sharp error bound is

$$0 < \pi - \pi_n < \pi^2 4^{n+2} \exp(-2\pi 4^n). \quad (29)$$

This improves by a factor of two on the error bound given in [14, top of pg. 171]. We defer the proof until §5.

Table 5 shows the error $\pi - \pi_n$ after n iterations of the Borwein quartic algorithm, and the ratio of the error $\pi - \pi_n$ to the upper bound (29).

n	$\pi - \pi_n$	$\frac{\pi - \pi_n}{\text{bound (29)}}$
0	2.273790912e-1	0.7710517124
1	7.376250956e-9	0.9602112619
2	5.472109145e-41	0.9900528160
3	2.308580715e-171	0.9975132040
4	1.110954934e-694	0.9993783010
5	9.244416653e-2790	0.9998445753
6	6.913088685e-11172	0.9999611438
7	3.376546688e-44702	0.9999902860
8	3.002256862e-178825	0.9999975715

Table 5: Approximation error in Algorithm BB4

At this point the reader may well ask “which of Algorithms GL, BB1, BB2 and BB4 is the fastest?”. The answer seems to depend on implementation details. All four algorithms involve the same number of square roots to obtain comparable accuracy (counting a fourth root in Algorithm BB4 as equivalent to two square roots, which is not necessarily correct⁷). Algorithm GL has the advantage that high-precision divisions are only required when generating the output (so the early divisions can be skipped if intermediate output is not required). The other three algo-

⁷ For example, one might compute $x^{1/4}$ using two *inverse* square roots, i.e. $(x^{-1/2})^{-1/2}$, which is possibly faster than two square roots, i.e. $(x^{1/2})^{1/2}$, see [22, §4.2.3].

rithms require at least one division per iteration. Borwein, Borwein and Bailey [15, pg. 202] say “[Algorithm BB4] is arguably the most efficient algorithm currently known for the extended precision calculation of π ”, and the times given in Bailey’s paper [4, pg. 289] confirm this (28 hours for Algorithm BB4 versus 40 hours for Algorithm BB1). However, Kanada [32], who extended Bailey’s computation, reached the opposite conclusion. His computation took 5 hours 57 minutes with Algorithm GL, and 7 hours 30 minutes with Algorithm BB4 (which was used for verification).

5 Equivalence of some algorithms for π

In the following, *doubling* an algorithm A means to construct an algorithm A^2 that outputs (x_0, x_2, x_4, \dots) if algorithm A outputs (x_0, x_1, x_2, \dots) . Replacing n by $2n$ in (20) and retaining only the most significant term, we see that an error bound for Algorithm GL1 doubled is

$$0 < \pi - a_{2n+1}^2/s_{2n} < \pi^2 4^{n+2} \exp(-2\pi 4^n).$$

It is suggestive that the right-hand side is the same as in the error bound (29) for the Borwein quartic algorithm after n iterations.

On closer inspection we find that the two algorithms (GL1 doubled and BB4) are *equivalent*, in the sense that they give *exactly* the same sequence of approximations to π . Symbolically,

$$\pi_n = a_{2n+1}^2/s_{2n}, \quad (30)$$

where a_n, s_n are as in Algorithm GL, and π_n is as in Algorithm BB4. This observation appears to be new – it is not stated explicitly in *Pi and the AGM* or elsewhere, so far as we know.⁸

Before proving the result, we give some empirical evidence for it, since that is how the result was discovered – in the spirit of “Experimental Mathematics”, as beloved by Jon Borwein. In Table 6, $n + 1$ is the number of square roots, and the second column is the error in the approximation given by Algorithm GL1 after n iterations, or by the Algorithm BB4 after $n/2$ iterations (n even). The error is the same for both algorithms (verified to 1000 decimal digits, not all shown). Using the definitions of the two algorithms, equality for the first line of the table ($n = 0$) follows from

$$a_1^2/s_0 = \pi_0 = \frac{3}{2} + \sqrt{2} = \pi - 0.227\dots$$

For the second line ($n = 2$) we have, with $t := 2^{-1/4}$,

$$a_3 = \frac{(t^2 + 2t + 1 + 2\sqrt{2t^3 + 2t})}{8} \quad \text{and} \quad s_2 = \frac{8t^3 - 4t^2 + 8t - 5}{16},$$

⁸ For example, the equivalence is not mentioned in [4], [15], [28], [29] or [32].

n	$\pi - a_{2n+1}^2/s_{2n}$ (for Algorithm GL1) or $\pi - \pi_n$ (for Algorithm BB4)
0	2.2737909121669818966095465906980480562749752399816e-1
2	7.3762509563132989512968071098827321760295030264154e-9
4	5.4721091456899418327485331789641785565936917028248e-41
6	2.3085807149343902668213207343869568303303472423996e-171
8	1.1109549335576998257002904117322306941479378545140e-694

Table 6: Approximation error for Algorithms GL1 doubled and BB4

so

$$\frac{a_3^2}{s_2} = \frac{(t^2 + 2t + 1 + 2\sqrt{2t^3 + 2t})^2}{4(8t^3 - 4t^2 + 8t - 5)}. \quad (31)$$

Also, from the definition of Algorithm BB4 we find, with

$$y_1 = \frac{1 - (12\sqrt{2} - 16)^{1/4}}{1 + (12\sqrt{2} - 16)^{1/4}},$$

that

$$\pi_1 = \frac{1}{(6 - 4\sqrt{2})(1 + y_1)^4 - 8y_1 - 8y_1^2 - 8y_1^3}. \quad (32)$$

It is not obvious that the algebraic numbers given by (31) and (32) are identical, but it can be verified that they both have minimal polynomial

$$\begin{aligned} P(x) := & 1 - 1635840576x - 343853312x^2 + 60576043008x^3 \\ & + 1865242664960x^4 - 16779556159488x^5 + 37529045696512x^6 \\ & - 29726424956928x^7 + 6181548457984x^8. \end{aligned}$$

Using Sturm sequences [44], it may be shown that $P(x)$ has two real roots, one in the interval $[0, 1]$, and the other in $[3, 4]$. A numerical computation shows that $|a_3^2/s_2 - \pi_1| < 1$, but both a_3^2/s_2 and π_1 are real roots of $P(x)$, so they must be equal.

Clearly this “brute force” approach does not generalise. To prove the equivalence of Algorithms BB4 and GL1, we first consider the equivalence of Algorithms BB2 and GL1.

Theorem 1. *Algorithm BB2 is equivalent to Algorithm GL1, in the sense that*

$$\widehat{\pi}_n = a_{n+1}^2/s_n,$$

where $\widehat{\pi}_n = 1/\alpha_n$ is as in Algorithm BB2, and a_{n+1}, s_n are as in Algorithm GL.

Proof. In the proof we take $n \geq 0$, $q = e^{-\pi}$, and assume that a_n, b_n, c_{n+1}, s_n are defined as in Algorithm GL, and $k_n, \alpha_n, \widehat{\pi}_n$ are as in Algorithm BB2.

Algorithm GL implements the recurrence

$$s_{n+1} = s_n - 2^n c_{n+1}^2, \quad (33)$$

whereas Algorithm BB2 implements the recurrence

$$\alpha_{n+1} = (1 + k_{n+1})^2 \alpha_n - 2^{n+2} k_{n+1}. \quad (34)$$

We show that the recurrences (33)–(34) are related. Noting the remark on subscripts following the statement of Algorithm BB2, we see that $k_n = c_{n+1}/a_{n+1}$, since both sides equal $\theta_2^2(q^{2^{n+1}})/\theta_3^2(q^{2^{n+1}})$. Thus

$$1 + k_{n+1} = a_{n+1}/a_{n+2}. \quad (35)$$

Define $\beta_n := a_{n+1}^2 \alpha_n$ and $\gamma_n := a_{n+2}^2 k_{n+1}$. Substituting (35) into (34) and clearing the fractions gives

$$\beta_{n+1} = \beta_n - 2^{n+2} \gamma_n. \quad (36)$$

Now

$$4\gamma_n = 4a_{n+2}^2 k_{n+1} = 4a_{n+2} c_{n+2} = \theta_2^4(q^{2^{n+2}})/\theta_3^4(q) = c_{n+1}^2,$$

so (36) is equivalent to

$$\beta_{n+1} = \beta_n - 2^n c_{n+1}^2. \quad (37)$$

This is essentially the same recurrence as (33). Also, $s_0 = 1/4$ and $\beta_0 = a_1^2 \alpha_0 = 1/4$, so $s_0 = \beta_0$. It follows that $s_n = \beta_n$ for all $n \geq 0$. Thus $s_n = a_{n+1}^2 \alpha_n$, and

$$\widehat{\pi}_n = 1/\alpha_n = a_{n+1}^2/s_n,$$

which completes the proof. \square

Corollary 1. *Algorithm BB4 is equivalent to Algorithm GL1 doubled, in the sense that*

$$\pi_n = a_{2n+1}^2/s_{2n},$$

where π_n is as in Algorithm BB4, and a_n, s_n are as in Algorithm GL.

Proof. The Borwein brothers noted [14, pg. 171] that Algorithm BB4 is equivalent to Algorithm BB2 doubled,⁹ i.e. $\pi_n = \widehat{\pi}_{2n}$. Thus, the result follows from Theorem 1. \square

Corollary 2. *For Algorithm BB4, the error bound (29) holds.*

Proof. In view of Corollary 1, the error bound (29) follows from (30) and the error bound (20) for Algorithm GL. \square

6 Some fast (but linear) algorithms for π

Let $(x)_n := x(x+1) \cdots (x+n-1)$ denote the *ascending factorial*. In Chapter 5 of *Pi and the AGM*, Jon and Peter Borwein discuss *Ramanujan-Sato* series such as

⁹ In fact, this is how Algorithm BB4 was discovered, by doubling Algorithm BB2 and then making some straightforward program optimisations.

$$\frac{1}{\pi} = 2^{3/2} \sum_{n=0}^{\infty} \frac{(\frac{1}{4})_n (\frac{1}{2})_n (\frac{3}{4})_n}{(n!)^3} \frac{(1103 + 26390n)}{99^{4n+2}}.$$

This is linearly convergent, with rate $1/99^4$, so adds nearly eight decimal digits per term, since $99^4 \approx 10^8$.

A more extreme example is the Chudnovsky series [23]

$$\frac{1}{\pi} = 12 \sum_{n=0}^{\infty} (-1)^n \frac{(6n)! (13591409 + 545140134n)}{(3n)! (n!)^3 640320^{3n+3/2}}, \quad (38)$$

which adds about 14 decimal digits per term.

Although such series converge only linearly, their convergence is so fast that they are competitive with higher-order algorithms such as Algorithm GL for computing highly accurate approximations to π . Which algorithm is the fastest in practice depends details of the implementation and on technological factors such as memory sizes and access times.

7 Fast algorithms for the elementary functions

In this section, we consider the *bit-complexity* of algorithms. The bit-complexity of an algorithm is the (worst case) number of single-bit operations required to complete the algorithm. For a fuller discussion, see Chapter 6 of *Pi and the AGM*. We are interested in asymptotic results, so are usually willing to ignore constant factors.

If all operations are performed to (approximately) the same precision, then it makes sense to count *operations* such as multiplications, divisions and square roots. Algorithms based on the AGM fall into this category.

If the precision of the operations varies widely, then bit-complexity is a more sensible measure of complexity. An example is Newton's method, which is self-correcting, so can be started with low precision. Another example is summing a series with rational terms, such as $e = \sum_{k=0}^{\infty} 1/k!$.

The bit-complexity of multiplying two n -bit numbers to obtain a $2n$ -bit product is denoted by $M(n)$. The classical algorithm shows that $M(n) = O(n^2)$, but various asymptotically faster algorithms exist. The best result so far, due to Harvey, van der Hoeven and Lecerf [30], is

$$M(n) = O\left(n \log n K^{\log^* n}\right)$$

with $K = 8$. Here the *iterated logarithm* function $\log^* n$ is defined by

$$\log^* n := \begin{cases} 0 & \text{if } n \leq 1; \\ 1 + \log^*(\log n) & \text{if } n > 1. \end{cases}$$

It is unbounded but grows *extremely* slowly as $n \rightarrow \infty$, e.g. slower than

$\log \log \cdots \log n$ [for any fixed number of logs].

Indeed, if the multiplication algorithm is implemented on a computer that fits in the observable universe and has components no smaller than atomic nuclei, then we can safely assume that $\log^* n$ is bounded by a moderate constant, and that multiplication has bit-complexity $O(n \log n)$.

We follow *Pi and the AGM* and assume that $M(n)$ is nondecreasing and satisfies the weak regularity condition

$$2M(n) \leq M(2n) \leq 4M(n).$$

Newton's method can be used to compute reciprocals and square roots with bit-complexity

$$O(M(n) + M(\lceil n/2 \rceil) + M(\lceil n/2^2 \rceil) + \cdots + M(1)) = O(M(n)).$$

It can be shown that the bit-complexities of squaring, multiplication, reciprocation, division, and root extraction are asymptotically the same, up to small constant factors [18]. All these operations have bit-complexity of order $M(n)$.

To compute π to n digits (binary or decimal) by the arctan formula (1), or to compute $1/\pi$ by the Chudnovsky series (38), we have to sum of order n terms. Using *divide and conquer*, also called *binary splitting* [18, 27],¹⁰ this can be done with bit-complexity

$$O(M(n) \log^2 n).$$

Suppose we compute π to n -digit accuracy using one of the quadratically convergent AGM algorithms. This requires $O(\log n)$ iterations, each of which has bit-complexity $O(M(n))$. Thus, the overall bit-complexity is

$$O(M(n) \log n).$$

This is (theoretically) better than series summation methods, the best of which have bit-complexity of order $M(n) \log^2 n$.

In practice, a method with bit-complexity of order $M(n) \log^2 n$ may be faster than a method with bit-complexity of order $M(n) \log n$ unless n is sufficiently large. This is one reason for the recent popularity of the Chudnovsky series (38) for high-precision computation of π , even though the AGM-based methods are theoretically (i.e. asymptotically) more efficient.

In §3, we mentioned Salamin's algorithm for computing $\log x$ for sufficiently large $x = 4/k$, i.e. sufficiently small k , using (9). We can evaluate $K'(k)/\pi$ using the AGM with $(a_0, b_0) = (1, k)$, and hence approximate $\log(4/k)$, assuming that π is precomputed. To compute $\log x$ to n -bit accuracy requires about $2 \log_2(n)$ AGM iterations, or $3 \log_2(n)$ iterations if we count the computation of π .

If x is not sufficiently large, we can use the identity $\log(x) = \log(2^p x) - p \log 2$, where p is a sufficiently large integer (but not too large or excessive cancellation will

¹⁰ Somewhat more general, but based on the same idea, is E. Karatsuba's *FEE method* [33].

occur). This assumes that $\log 2$ is precomputed, and that the precision is increased to compensate for cancellation.

To obtain a small relative error when x is close to 1, say $|x - 1| < 2^{-n/\log n}$, it is better to use the Taylor series for $\log(1 + z)$, with $z = x - 1$. The Taylor series computation can be accelerated by “splitting”, see [22, §4.4.3] and [43],

The $O(k^2)$ error term in the expression (9) can be written explicitly using hypergeometric series, see [14, (1.3.10)]. This gives one way of improving the accuracy of the approximation $K'(k)$ to $\log(4/k)$. We give an alternative using theta functions, for which the series converge faster than the hypergeometric series (which converge only linearly). The result (39) follows from several identities given in §2. We collect them here for convenience:

$$\begin{aligned}\log(1/q) &= \pi K'(k)/K(k), \\ k &= \theta_2^2(q)/\theta_3^2(q), \\ K(k) &= (\pi/2)\theta_3^2(q), \\ K'(k) &= (\pi/2)/\text{AGM}(1, k).\end{aligned}$$

Putting these pieces together gives the elegant result of Sasaki and Kanada [42]

$$\log(1/q) = \frac{\pi}{\text{AGM}(\theta_2^2(q), \theta_3^2(q))}. \quad (39)$$

In (39) we can replace q by q^4 to avoid fractional powers of q in the expansion of $\theta_2(q)$, obtaining an exact formula for all $q \in (0, 1)$:

$$\log(1/q) = \frac{\pi/4}{\text{AGM}(\theta_2^2(q^4), \theta_3^2(q^4))}. \quad (40)$$

As in Salamin’s algorithm, we have to ensure that $x := 1/q$ is sufficiently large, but now there is a trade-off between increasing x or taking more terms in the series defining the theta functions. For example, to attain n -bit accuracy, if $x > 2^{n/36}$, we can use $\theta_2(q^4) = 2(q + q^9 + q^{25} + O(q^{49}))$ and $\theta_3(q^4) = 1 + 2(q^4 + q^{16} + O(q^{36}))$. This saves about four AGM iterations, compared to Salamin’s algorithm. We remark that a result similar to (39) and (40) is given in (7.2.5) of *Pi and the AGM*, but with an unfortunate typo (a reciprocal is missing).

So far we have assumed that the initial values a_0, b_0 in the AGM iteration are real and positive. There is no difficulty in extending the results that we have used to complex a_0, b_0 , provided that they are nonzero and a_0/b_0 is not both real and negative. For simplicity, we assume that $a_0, b_0 \in \mathcal{H} = \{z \mid \Re(z) > 0\}$.

In the AGM iteration (and in the definition of the geometric mean) there is an ambiguity of sign. We always choose the square root with positive real part. Thus the iterates a_n, b_n are uniquely defined and remain in the right half-plane \mathcal{H} .

When using (40), we may need to apply a rotation to q , say by a multiple of $\pi/3$, in order to ensure that the starting values $(\theta_2^2(q^4), \theta_3^2(q^4))$ for the AGM lie in \mathcal{H} .¹¹

For $z \in \mathbb{C} \setminus \{0\}$, $\log(z) = \log(|z|) + i \arg(z)$, provided we use the principal values of the logarithms. Thus, if $x \in \mathbb{R}$, we can use the complex AGM to compute

$$\arctan(x) = \Im(\log(1 + ix)).$$

$\arcsin(x)$, $\arccos(x)$ etc can be computed via \arctan using elementary trigonometric identities such as $\arccos(x) = \arctan(\sqrt{1-x^2}/x)$.

Since we can compute \log , \arctan , \arccos , \arcsin , we can compute \exp , \tan , \cos , \sin (in suitably restricted domains) using Newton's method. The trigonometric functions can also be computed via the complex exponential. Similarly for the hyperbolic functions \cosh , \sinh , \tanh and their inverse functions.

Although computing the elementary functions via the complex AGM is conceptually straightforward, it introduces the overhead of complex arithmetic. It is possible to avoid complex arithmetic by the use of Landen transformations (which transform incomplete elliptic integrals). See exercise 7.3.2 of *Pi and the AGM* for an outline of this approach, and [19] for more details.

Whichever approach is used, the bit-complexity of computing n -bit approximations to any of the elementary functions (\log , \exp , \arctan , \sin , \cos , \tan , etc) in a given compact set $A \subset \mathbb{C}$ that excludes singularities of the relevant function is $O(M(n) \log n)$. Here “ n -bit approximation” means with absolute error bounded by 2^{-n} . We could require relative error bounded by 2^{-n} , but the proof would depend on a Diophantine approximation result such as Mahler's well-known result on approximation of π by rationals [37], because of the difficulty of guaranteeing a small relative error in the neighbourhood of a zero of the function.¹²

Certain non-elementary functions can be computed with bit-complexity $O(M(n) \log n)$ via the AGM. For example, we mention complete and incomplete elliptic integrals, elliptic functions, and the Jacobi theta functions $\theta_2(q)$, $\theta_3(q)$, $\theta_4(q)$. Functions that appear *not* to be in this class of “easily computable” functions include the Gamma function $\Gamma(z)$ and the Riemann zeta function $\zeta(s)$.

Algebraic functions can be computed with bit-complexity $O(M(n))$, see [14, Thm. 6.4]. It is plausible to conjecture that no elementary transcendental functions can be computed with bit-complexity $O(M(n))$. However, as usual in complexity theory, nontrivial lower bounds are difficult to prove.

¹¹ Alternatively, we could drop the simplifying assumption that $a_0, b_0 \in \mathcal{H}$ and use the “right choice” of Cox [24, pg. 284] to implement the AGM correctly.

¹² Mahler's result is sufficient for the usual elementary functions, whose zeros are rational multiples of π , but it is not applicable to the problem of computing combinations of these functions, e.g. $\exp(\sin x) + \cos(\log x)$, with small relative accuracy. In general, we do not know enough about the rational approximation of the zeros of such functions to guarantee a small *relative* error. However, the result that we stated for computing elementary functions with a small *absolute* error extends to finite combinations of elementary functions under the operations of addition, multiplication, composition, etc. Indeed, the set of *elementary functions* is usually considered to include such finite combinations, although precise definitions vary. See, for example, §7.3 of *Pi and the AGM*, Knopp [34, pp. 96–98], Liouville [36], Ritt [40], and Watson [45, pg. 111].

Acknowledgements I am grateful to Jon and Peter Borwein for becoming sufficiently interested in this subject to write their book *Pi and the AGM* only a few years after the publication of [17, 18, 19, 41]. Reading a copy of *Pi and the AGM* was my first introduction to the Borwein brothers, and was the start of my realisation that we shared many common interests, despite living in different hemispheres. Much later, after Jon moved to Newcastle, I followed him, bringing our common interests closer together, and benefitting from frequent interaction with him.

The author was supported in part by Australian Research Council grant DP140101417 (Jon Borwein was the Principal Investigator on this grant, up to the time of his death).

References

1. M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions*, Dover, New York, 1965. Online version at <http://people.math.sfu.ca/~cbm/aands/>, accessed 19 February, 2018.
2. G. E. Andrews, Book review: *Pi and the AGM* . . . , *Bulletin (NS) AMS* **22** (1990), 198–201.
3. R. Askey, Book review: *Pi and the AGM*, *Amer. Math. Monthly* **95** (1988), 895–897.
4. D. H. Bailey, The computation of π to 29,360,000 decimal digits using Borweins' quartically convergent algorithm, *Math. Comp.* **50** (1988), 283–296.
5. D. H. Bailey, A collection of mathematical formulas involving π , Feb. 6, 2018, <http://www.davidhbailey.com/dhbpapers/pi-formulas.pdf>, accessed 19 February, 2018.
6. D. H. Bailey and J. M. Borwein, *Pi: The Next Generation*, Springer, 2016.
7. N. D. Baruah, B. C. Berndt and H. H. Chan, Ramanujan's series for $1/\pi$: a survey, *Amer. Math. Monthly* **116** (2009), 567–587.
8. M. Beeler, R. W. Gosper and R. Schroepel, *HAKMEM*, AI Memo 239, MIT AI Lab, Feb. 1972. (Item 143 by E. Salamin.)
9. B. C. Berndt, Book review: *Pi and the AGM* . . . , *Math. Comput.* **50** (1988), 352–354.
10. J. M. Borwein, The life of pi: from Archimedes to Eniac and beyond, prepared for Berggren Festschrift, 19 June 2012, <https://www.carma.newcastle.edu.au/jon/pi-2012.pdf>, accessed 19 February, 2018.
11. J. M. Borwein, Lectures and Presentations, <https://www.carma.newcastle.edu.au/jon/index-talks.shtml>, accessed 19 February, 2018.
12. J. M. Borwein and P. B. Borwein, The arithmetic-geometric mean and fast computation of elementary functions, *SIAM Review* **26** (1984), 351–365.
13. J. M. Borwein and P. B. Borwein, More quadratically convergent algorithms for π , *Math. Comput.* **46** (1986), 247–253.
14. J. M. Borwein and P. B. Borwein, *Pi and the AGM*, Monographies et Études de la Société Mathématique du Canada, John Wiley & Sons, Toronto, 1987.
15. J. M. Borwein, P. B. Borwein and D. H. Bailey, Ramanujan, modular equations, and approximations to pi or how to compute one billion digits of pi, *Amer. Math. Monthly* **96** (1989), 201–219.
16. R. P. Brent, Some efficient algorithms for solving systems of nonlinear equations, *SIAM J. Numer. Anal.* **10** (1973), 327–344.
17. R. P. Brent, Multiple-precision zero-finding methods and the complexity of elementary function evaluation, in *Analytic Computational Complexity* (edited by J. F. Traub), Academic Press, New York, 1975, 151–176.
18. R. P. Brent, The complexity of multiple-precision arithmetic, in *The Complexity of Computational Problem Solving* (R. S. Anderssen and R. P. Brent, eds.), Univ. of Qld. Press, Brisbane, 1976, 126–165.
19. R. P. Brent, Fast multiple-precision evaluation of elementary functions, *J. ACM* **23** (1976), 242–251.

20. R. P. Brent, Old and new algorithms for π , *Notices of the AMS* **60** (Jan. 2013), 7.
21. R. P. Brent, *Jonathan Borwein, Pi and the AGM*, keynote talk at the Jonathan Borwein Commemorative Conference, Newcastle, NSW, 26 Sept. 2017, <http://maths-people.anu.edu.au/~brent/talks.html>, accessed 19 February, 2018.
22. R. P. Brent and P. Zimmermann, *Modern Computer Arithmetic*, Cambridge University Press, 2010.
23. D. V. Chudnovsky and G. V. Chudnovsky, The computation of classical constants, *Proc. Nat. Acad. Sci. USA* **88**(21), 8178–8182.
24. D. A. Cox, The arithmetic-geometric mean of Gauss, *L'Enseignement Mathématique* **30** (1984), 275–330.
25. C. F. Gauss, unpublished notebook entry of May 1809, reproduced in J. Arndt and C. Haenel, *Pi: Algorithmen, Computer, Arithmetik*, Springer, Berlin, 1998, Ch. 7, pg. 99.
26. C. F. Gauss, *Carl Friedrich Gauss Werke*, Bd. 3, Göttingen, 1876, 362–403.
27. X. Gourdon and P. Sebah, *Binary splitting method*, 2001, <http://numbers.computation.free.fr/Constants/Algorithms/splitting.html>, accessed 19 February, 2018.
28. J. Guillera, Easy proofs of some Borwein algorithms for π , *Amer. Math. Monthly* **115** (2008), 850–854.
29. J. Guillera, New proofs of Borwein-type algorithms for Pi, *Integral Transforms and Special Functions* **27** (2016), 775–782.
30. D. Harvey, J. van der Hoeven and G. Lecerf, Even faster integer multiplication, *J. Complexity* **36** (2016), 1–30.
31. C. G. J. Jacobi, *Fundamenta Nova Theoriae Functionum Ellipticarum*, Königsberg, 1829. Reprinted in *Gesammelte Mathematische Werke*, Bd. 1, 255–263.
32. Y. Kanada, Vectorization of multiple-precision arithmetic program and 201,326,000 decimal digits of pi calculation, *Supercomputing* 88, IEEE, 1988, 117–128.
33. E. A. Karatsuba, Fast evaluations of transcendental functions, *Probl. Peredachi Informat.* **27**, 4 (1991). Also https://en.wikipedia.org/wiki/FEE_method, accessed 19 February, 2018.
34. K. Knopp, *The Elementary Functions*, §23 in *Theory of Functions Parts I and II*, Dover, New York, 1996, 96–98.
35. A. M. Legendre, *Exercices de Calcul Integral*, Vol. 1, Paris, 1811, pg. 61.
36. J. Liouville, Sur la classification des Transcendantes et sur l'impossibilité d'exprimer les racines des certaines équations en fonction finie explicite des coefficients. Part 1, *J. Math. Pure Appl.* **2** (1837), 56–105. Also Part 2, *ibid* **3** (1838), 523–547.
37. K. Mahler, On the approximation of π , *Proc. Kon. Nederlandsche Akad. v. Wetenschappen Ser. A* **56** (1953), 30–42 = *Indag. Math.* **15** (1953), 30–42. Also <https://carma.newcastle.edu.au/mahler/docs/119.pdf>, accessed 19 February, 2018.
38. A. M. Ostrowski, *Solution of Equations and Systems of Equations*, Academic Press, New York, 1960.
39. S. Ramanujan, Modular equations and approximations to pi, *Quart. J. Math. (Oxford)* **45** (1914), 350–372.
40. J. F. Ritt, *Integration in Finite Terms*, Columbia Univ. Press, New York, 1948.
41. E. Salamin, Computation of π using arithmetic-geometric mean, *Math. Comp.* **30** (1976), 565–570.
42. T. Sasaki and Y. Kanada, Practically fast multiple-precision evaluation of $\log(x)$, *J. Inf. Process.* **5** (1982), 247–250.
43. D. M. Smith, Efficient multiple-precision evaluation of elementary functions, *Math. Comp.* **52** (1989), 131–134.
44. J. C. F. Sturm, Mémoire sur la résolution des équations numériques, *Bulletin des Sciences de Férussac* **11** (1829), 419–425.
45. G. N. Watson, *A Treatise on the Theory of Bessel Functions*, second edn., Cambridge, 1966.
46. E. T. Whittaker and G. N. Watson, *A Course of Modern Analysis*, third edition, Cambridge, 1920. Also <http://archive.org/details/cu31924001549660>, Accessed 19 February, 2018.
47. J. Wimp, Pi and the AGM . . . , review in *SIAM Review* **30** (1988), 530–533.