

Verifying an Efficient Algorithm for Computing Bernoulli Numbers

Manuel Eberl  

University of Innsbruck, Austria

Peter Lammich  

University of Twente, The Netherlands

Abstract

The Bernoulli numbers B_k are a sequence of rational numbers that is ubiquitous in mathematics, but difficult to compute efficiently (compared to e.g. approximating π).

In 2008, Harvey gave the currently fastest known practical way for computing them: his algorithm computes $B_k \bmod p$ in time $O(p \log^{1+o(1)} p)$. By doing this for $O(k)$ many small primes p in parallel and then combining the results with the Chinese Remainder Theorem, one recovers the value of B_k as a rational number in $O(k^2 \log^{2+o(1)} k)$ time. One advantage of this approach is that the expensive part of the algorithm is highly parallelisable and has very low memory requirements. This algorithm still holds the world record with its computation of B_{10^8} .

We give a fully verified efficient LLVM implementation of this algorithm. This was achieved by formalising the necessary mathematical background theory in Isabelle/HOL, proving an abstract version of the algorithm correct, and refining this abstract version down to LLVM using Lammich's *Isabelle-LLVM* framework, including many low-level optimisations. The performance of the resulting LLVM code is comparable with Harvey's original unverified and hand-optimised C++ implementation.

2012 ACM Subject Classification Theory of computation \rightarrow Logic and verification

Keywords and phrases Bernoulli numbers, LLVM, verification, Isabelle, Chinese remainder theorem, modular arithmetic, Montgomery arithmetic

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

Acknowledgements The computational results presented here have been achieved (in part) using the LEO HPC infrastructure of the University of Innsbruck.

1 Introduction

The Bernoulli numbers B_k are a sequence of rational numbers that are ubiquitous in mathematics: they have connections to, among other things: the closed-form expression for $\sum_{i=1}^n i^k$ (Faulhaber's formula), regular primes and cyclotomic fields, the combinatorics of alternating permutations, the Riemann zeta function, the Euler–Maclaurin summation formula, and the Maclaurin series of e.g. the tangent and cosecant functions.

Algorithmically, they are quite difficult to compute efficiently: Bernoulli himself only computed them up to $B_{10} = 5/66$ in the 17th century. Euler extended this up to $B_{30} = 8615841276005/14322$ in 1748, Adams up to B_{124} in 1877 [1], and Lehmer up to B_{220} in 1936 [21] – all by hand.

Interestingly, it is often claimed that the first non-trivial computer program ever written was Lovelace's *Note G*, published in 1843. This was an algorithm designed for Babbage's *Analytical Engine* (which was sadly never built) to compute Bernoulli numbers.¹ [8]

The first actual use of a digital computer (that we are aware of) in computing Bernoulli numbers (and the first big step forward after Lehmer in 1936) was when Knuth and Buckholtz

¹ Accounts differ on whether it was Lovelace or Babbage himself who authored the program, but the prevailing view seems to be that it was Lovelace.

42 gave a table of the values up to B_{836} in 1967 [15] (the numerator of B_{836} already has 1421
 43 decimal digits). The first large jump to higher indices happened in 1996, when Fee and
 44 Plouffe [9] computed the single value $B_{750,000}$ and a few others up to $B_{5,000,000}$ in the following
 45 years, using an entirely different method based on approximating the Riemann zeta function.
 46 Unlike previous methods, this method allowed computing a value B_k without knowing the
 47 preceding values B_0, \dots, B_{k-1} .

48 In 2010, Harvey [13] computed the value of B_{10^8} using an entirely different method: He
 49 gave a fast algorithm to compute B_k modulo a prime p . This can then be run in parallel
 50 for many different primes, followed by an invocation of the Chinese Remainder Theorem, to
 51 reconstruct the value of B_k as a rational number. While this approach is not asymptotically
 52 faster than the one based on the zeta function, it is much faster in practice, as Harvey
 53 demonstrated. To this day, Harvey’s algorithm is still the state of the art for practical
 54 computation of a single Bernoulli number, and his record of B_{10^8} still stands. His algorithm
 55 takes time $O(k^2 \log^{2+o(1)} k)$ to compute B_k .²

56 In this paper, we verify Harvey’s algorithm in Isabelle/HOL. Using the Isabelle Refinement
 57 Framework [20] and its parallel LLVM back end [19], our verification spans from a definition
 58 of Bernoulli numbers down to efficient LLVM code with performance comparable to Harvey’s
 59 original algorithm. As a side product, we also verified several other interesting numerical
 60 algorithms.

61 2 Mathematical Background

62 The Bernoulli numbers B_k are defined as the coefficients of the exponential generating
 63 function $B(z) = z/(\exp(z) - 1)$ or, following an alternative convention, $z/(1 - \exp(-z))$.
 64 The only difference between the two conventions is that the first one yields $B_1 = -\frac{1}{2}$, and
 65 the second one yields $B_1 = \frac{1}{2}$. We shall adopt the first convention. Furthermore, we write
 66 N_k and D_k for the numerator and denominator of B_k , respectively. We make use of the
 67 Isabelle/HOL formalisation of Bernoulli numbers in the Archive of Formal Proofs (AFP) by
 68 Bulwahn and Eberl [4], which already contains many of the results on Bernoulli numbers
 69 that we shall use.

70 ► **Remark 1 (Notation).** Both in this presentation and in our Isabelle proofs, we use the
 71 notation $x \equiv_n y$. If x, y are integers, the meaning is that $x \bmod n = y \bmod n$ or, equivalently,
 72 $n \mid (x - y)$. The corresponding notation in Isabelle/HOL is $[x = y] \pmod n$.

73 We extend this definition to rational numbers in the natural way: We define $\frac{a}{b} \bmod n =$
 74 $ab^{-1} \bmod n$, where b^{-1} denotes a modular inverse of b modulo n . Note that this is only
 75 well-defined if the denominators are coprime to n . The Isabelle notation we introduced
 76 for this is $x \text{ qmod } n$ for the infix operator ‘mod’ and $[x = y] \text{ (qmod } n)$ for the congruence
 77 relation.

78 Also note that, as is common in number theory, any sum and product with an index
 79 variable p is meant to be interpreted with p being restricted to the prime numbers.

80 The sequence of Bernoulli numbers shows an obvious pattern: if $k > 1$ then $B_k = 0$ if k
 81 is odd, $B_k > 0$ if $k \equiv_4 2$, and $B_k < 0$ if $k \equiv_4 0$. The vanishing of B_k for odd $k > 1$ is easily
 82 proved from the formal power series given above. The result about the sign of B_k for even
 83 $k > 1$ is obtained from a well-known connection between B_k and the Riemann zeta function:

² For comparison, π has been approximated to trillions of digits and is much easier to compute ($O(n \log^3 n)$ time for n digits with Chudnovsky’s formula and binary splitting). [24]

84 ▶ **Theorem 2** (The connection between B_k and ζ). *If $k > 0$, we have:*

$$85 \quad \zeta(2k) = \sum_{n \geq 1} n^{-2k} = \frac{(-1)^{k+1} (2\pi)^{2k} B_{2k}}{2(2k)!}$$

86 **Proof.** We apply the Residue Theorem to the integral $\oint_{R_m} B(z)/z^{2k+1} dz$. Here, R_m is a
87 rectangle of width $2m$ and height $2(2m+1)\pi$ centred around the origin. Letting $m \rightarrow \infty$
88 establishes our claim. ◀

89 The next important ingredient gives us an easy way to compute the denominator of B_k :

90 ▶ **Theorem 3** (Von Staudt–Clausen theorem). *If $k \geq 2$ is even, then $D_k = \prod_{(p-1)|k} p$.*

91 We omit the proofs of this fact and of the following ones from this presentation for reasons
92 of space. All results up to this point are part of the AFP entry by Bulwahn and Eberl [4].

93 The core of Harvey's algorithm are two closely related congruences involving Bernoulli
94 numbers modulo a prime or a prime power:

95 ▶ **Theorem 4** (Voronoi's congruence). *Let $k \geq 2$ be even and let $n > 0$ and a be coprime
96 integers. Then:*

$$97 \quad (a^k - 1)N_k \equiv_n ka^{k-1}D_k \sum_{1 \leq m < n} m^{k-1} \left\lfloor \frac{ma}{n} \right\rfloor$$

98 ▶ **Theorem 5** (Kummer's congruence). *Let p be a prime number and $e \geq 0$ and $k, k' \geq e + 1$
99 be integers with k, k' even and $k \equiv_{p^{e-1}(p-1)} k'$ and $(p-1) \nmid k$. Then $B_k/k \equiv_{p^e} B_{k'}/k'$.*

100 **Proof.** The Isabelle proofs closely follow Cohen [5, Prop. 9.5.20, Cor. 9.5.25] and are relatively
101 straightforward. ◀

102 These congruences were not previously available in the AFP and were formalised in the
103 course of this project [7].

104 By setting $e := 1$ and $k' := k \bmod (p-1)$ in Kummer's congruence, we obtain the
105 range-reduction congruence $B_k \equiv_p k/k' B_{k'}$. Thus we can w.l.o.g. assume that $2 \leq k \leq p-3$.

106 Voronoi's congruence gives us a closed-form expression for computing $B_k \bmod p$. Harvey
107 tweaks this congruence to obtain the following:

108 ▶ **Theorem 6** (Voronoi's congruence, Harvey's first version). *Let $p \geq 5$ be prime and $2 \leq k \leq$
109 $p-3$ and $1 < c < p$ with $c^k \not\equiv_p 1$. Then:*

$$110 \quad B_k \equiv_p \frac{k}{1-c^k} \sum_{1 \leq m < p} m^{k-1} h(m) \quad \text{where } h(m) = \frac{m - c((mc^{-1}) \bmod p)}{p} + \frac{c-1}{2} \in \mathbb{Q}$$

111 Here, c^{-1} denotes a modular inverse of c modulo p . Note that $h(m)$ is either an integer or a
112 half-integer, depending on whether c is even or odd.

113 We will later pick a convenient value for the parameter c in this theorem.

114 **3** Computing N_k Modulo a Prime

115 The main part of Harvey’s algorithm is to gather ‘modular information’, i.e. to compute
 116 $B_k \bmod p$ for many different primes p . Note that since $p \mid D_k$ if $(p-1) \mid k$, the term $B_k \bmod p$
 117 is actually undefined if $(p-1) \mid k$, which is why Harvey does not use such primes.

118 We on the other hand decided to instead compute $N_k \bmod p$, which is well-defined for all
 119 primes. Since $D_k \bmod p$ is easy to compute via von Staudt–Clausen, it is easy to convert
 120 between $N_k \bmod p$ and $B_k \bmod p$ if the latter is well-defined.

121 Depending on p and k , we use one of three different methods to compute $N_k \bmod p$:

122 **Case 1:** If $(p-1) \nmid k$, we use a version of Theorem 6 to compute $B_k \bmod p$ and then multiply
 123 the result with $D_k \bmod p$ to obtain $N_k \bmod p$. To avoid big-integer arithmetic, we do
 124 the latter part by simply multiplying with $q \bmod p$ for every prime q with $(q-1) \mid k$.

125 There are two sub-cases for how we compute $B_k \bmod p$:

126 **Case 1.1** If $2^k \not\equiv_p 1$, we use a highly optimised version of Theorem 6 which we refer to
 127 as `harvey_fastsum` in our Isabelle formalisation.

128 **Case 1.2** If $2^k \equiv_p 1$, we fall back to Harvey’s *Algorithm 1*, which is a much less optimised
 129 version of Theorem 6. In our Isabelle formalisation, we refer to this as `harvey_slowsum`.

130 **Case 2** If $(p-1) \mid k$, we recall that $B_k \bmod p$ is not well-defined. However, it is easy to
 131 show that in this case we have $N_k \equiv_p -D_k/p$, which is easy to compute.

132 Since Case 1.1 is by far the most frequent one and also the most optimised one, we dedicate
 133 the remainder of this section to it. For Case 1.2, we refer the reader to Harvey’s paper.

134 Harvey derives the congruence behind Case 1.1 from Theorem 6 by choosing $c = \frac{1}{2} \bmod p$.
 135 He then picks a generator g of $\mathbb{Z}/p\mathbb{Z}$ (also referred to as a primitive root modulo p) and,
 136 with some elementary algebra and clever reindexing, derives the following:

137 **► Theorem 7 (The `harvey_fastsum` congruence).** *Let $p \geq 5$ be prime with $2^k \not\equiv_p 1$ and*
 138 *$2 \leq k \leq p-3$ and g a generator of $\mathbb{Z}/p\mathbb{Z}$. Let $n = \text{ord}_p(2)/2$ if $\text{ord}_p(2)$ is even and*
 139 *$n = \text{ord}_p(2)$ otherwise. Let $m = (p-1)/(2n)$. Then:*

$$140 \quad B_k \equiv_p \frac{k}{2(2^{-k}-1)} \sum_{0 \leq i < m} g^{i(k-1)} \sum_{0 \leq j < n} (2^{k-1})^j \sigma(p, 2^j g^i) \quad \text{where } \sigma(p, x) = \text{sgn}(p-2(x \bmod p))$$

141 Empirically, m is small for most primes.³ Thus, just like Harvey, we focus on optimising the
 142 inner sum. We will process the inner sum in chunks of w' blocks of size w . Currently, we
 143 pick $w = w' = 8$ for hardware architecture reasons that will become clear later.

144 Concretely, generalise 2^{k-1} to x and g^i to s in the inner sum in Theorem 7 and let
 145 $n_1 = \lfloor n/(ww') \rfloor$ and $n_2 = n \bmod (ww')$. Then:

$$146 \quad \sum_{0 \leq j < n} x^j \sigma(p, 2^j s) = \left(\sum_{0 \leq j < n_1 ww'} x^j \sigma(p, 2^j s) \right) + x^{n_1 ww'} \left(\sum_{0 \leq j < n_2} x^j \sigma(p, 2^{j+n_1 ww'} s) \right)$$

147 That is, we split the sum into a ‘chunk-aligned’ part, and a remainder.

³ The mean value of m for the first 10^6 odd primes is approximately 9.80. Less than 3% of these primes have $m \geq 30$.

148 3.1 The Inner Sum in General

149 For presentation purposes, we start by describing the algorithm for the remainder, which
 150 works for arbitrary n . The algorithm for chunk-aligned n uses the same implementation
 151 techniques with the addition of a tabulation optimisation, which we describe in Section 3.2.

152 3.1.1 Abstract Algorithm

153 We compute the sum $\sum_{0 \leq j < n} x^j \sigma(p, 2^j s)$ for arbitrary n with a simple loop with two local
 154 variables z and s' that hold the current values of $x^j \bmod p$ and $2^j s \bmod p$, respectively.

155 We have $\sigma(p, 2^j s) = \sigma(p, 2^j s \bmod p) = \sigma(p, s')$ and also $\sigma(p, s') = -1$ if $p \leq 2s'$ and
 156 $\sigma(p, s') = 1$ otherwise. This leads us to the following simple algorithm:

```
157
158 1 harvey_restsum1 p x s n = doN {
159 2   (_,_,acc) ← for 0 n (λj (z,s',acc). doN {
160 3     assert (s' ∈ {0..

```

165 This algorithm is phrased in the nondeterminism-error monad of the Isabelle Refinement
 166 Framework (IRF) [20]. It iterates over the values $j = 0, 1, \dots, n - 1$, maintaining the state
 167 (z, s', acc) with the invariants $z = x^j$ and $s' = 2^j s \bmod p$ and $acc \equiv_p \sum_{0 \leq j' < j} x^{j'} \sigma(p, 2^{j'} s)$.

168 Given the ideas above, we can prove the algorithm correct:

```
169
170 1 lemma harvey_restsum1_correct:
171 2   assumes "s ∈ {0..

```

175 In words: assuming that $0 \leq s < p$ and that p is odd, our algorithm refines the program
 176 that just returns the desired sum.

177 If m and m' are programs, the refinement relation $m \leq m'$ means that every possible
 178 result of program m is also a possible result of program m' , or m' fails. Moreover, m can
 179 only fail if m' fails. In the above lemma, m' is a deterministic program that does not fail,
 180 thus our algorithm does not fail and has either no results or the desired result. The case of
 181 having no results will only be excluded later, when we subsequently refine our algorithm to
 182 Isabelle-LLVM, which provably cannot have an empty set of results.

183 An assertion fails if the asserted predicate does not hold. In the context of stepwise
 184 refinement, assertions are a valuable tool to organise proof obligations: during the proof of
 185 the algorithm, we know that $s' < p$ (it follows from the invariant $s' = 2^j s \bmod p$). Making
 186 this knowledge explicit as an assertion allows us to use it later on, when we further refine
 187 the algorithm: since a failed algorithm is refined by any program, we can assume that the
 188 algorithm to be refined does not fail, and thus that the assertions hold. Also, as in unverified
 189 programming, assertions are an engineering tool, making flawed proof attempts fail early
 190 and helping us analyse the problem. For the sake of readability, we will often elide assertions
 191 from presented listings.

192 For this paper, we will elide most actual Isabelle proofs and rather describe their ideas.
 193 The Isabelle listings shown are slightly edited for presentation purposes. Lemma and function
 194 names coincide with the actual formalisation.

23:6 Verifying an Efficient Algorithm for Computing Bernoulli Numbers

```

1 harvey_restsum2 W p' p x s n = doN {
2   z ← to_mont1 W p p' 1;
3   (_,_,acc) ← for 0 n (λ_ (z,s',acc)). doN {
4     pl ← cast_u.op (2*W) p; s' ← cast_u.op (2*W) s'; s' ← mul_uuu.op (2*W) 2 s';
5     z' ← mont_times1 W p p' z x;
6     if s' < pl then doN {
7       acc ← mont_add_relaxed1 W acc z; s' ← cast_u.op W s'; return (z',s',acc)
8     } else doN {
9       acc ← mont_diff_relaxed1 W p acc z;
10      s' ← sub_uuu.op (2*W) s' pl; s' ← cast_u.op W s';
11      return (z',s',acc) }
12   } (z,s,mont_zero_relaxed1 W);
13   return acc }

```

■ **Figure 1** Computing the inner sum using Montgomery form

195 3.1.2 Introducing Montgomery Form

196 In a next step, we use the Montgomery arithmetic (also known as REDC arithmetic) [22] for
197 the accumulator and x^j . We also insert annotations for type casting. The algorithm is shown
198 in Fig. 1. Here, W is the bit width for numeric operations (in our case 32). The accumulator
199 is maintained as a double-width word that will not overflow during the loop. Thus, we can
200 defer the reduction modulo p until after the loop.

201 Moreover, the code contains annotations for type casts, bit widths, and the types of
202 operations. For example, the operation `mul_uuu.op (2*W) a b` multiplies its unsigned double-
203 width ($2W$) arguments a and b , asserting that there is no overflow. While these annotations
204 are not strictly necessary, they make the next refinement step simpler (cf. Sec. 3.1.3).

205 Some operations for the Montgomery form also require extra parameters apart from the
206 bit width W . For example, the multiplication requires both the modulus p and p' , its inverse
207 modulo 2^W . The inverse p' is pre-computed once and passed as an extra parameter to the
208 function. We show that this function is a refinement of the above `harvey_restsum1`:

```

209
210 1 lemma harvey_restsum2_refine:
211 2   assumes "mont_ctxt_refine' W p p' ctxt" "n < 2^W" "(xi,x) ∈ mont_rel ctxt"
212 3   shows "harvey_restsum2 W p' p xi s n
213 4     ≤ ↓(mont_rel_relaxed ctxt (n+1)) (harvey_restsum1 p x s n)"

```

215 The first assumption states that `ctxt` is a Montgomery refinement context⁴ for bit width W ,
216 modulus p and its inverse p' . The lemma then states that if n is in bounds, and xi is the
217 Montgomery form of x , then `harvey_restsum2 W p' p xi s n` refines `harvey_restsum1 p x s n`
218 wrt. the relation `mont_rel_relaxed ctxt (n+1)`. This relation indicates that the left value
219 modulo p is the Montgomery form of the right value, and that the left value is less than
220 $(n+1)p$. The counter $(n+1)$ is used to keep track of an upper bound for the accumulator,
221 while we keep adding values $< p$ to it.

222 Note that we can easily prove the non-overflow assertions on this level. The assertion
223 $s' < p$ from the previous refinement level helps here, as it allows us to assume $s' < p$.

⁴ The formalisation of Montgomery form uses these contexts, while our algorithms use explicit parameters. This are independent design choices, which slightly clash here.

3.1.3 Refinement to LLVM

The last step uses the Sepref [17] tool to refine to (an Isabelle model of) LLVM code.

```

226
227 1 sepref_def harvey_restsum_impl is "harvey_restsum2 32"
228 2   :: "u32A * u32A * u32A * u32A * u64A → u64A"
229 3   unfolding harvey_restsum2_def for_by_while
230 4   apply (annot_unat_const "TYPE(64)", annot_uint_const "TYPE(32)")
231 5   by sepref
232

```

The already proved bounds annotations make this proof easy: after unfolding the definition and unfolding the for loop into a while loop (which Sepref can process), we annotate the number literals with their respective bit-width, and the invoke Sepref. This generates LLVM code and a refinement lemma, stating that the parameters p , p' , xi , s are implemented by unsigned 32-bit words (u32A), the parameter n is implemented by an unsigned 64-bit word (u64A), and the result is returned as an unsigned 64-bit word.

3.2 Further Optimisation of the Inner Sum

For the aligned part of the summation, we apply another crucial optimisation that is also described by Harvey. Instead of iterating through the sum in single steps, we iterate in chunks of size ww' . Note that $\sigma(p, 2^j s)$ being 1 or -1 is equivalent to the j th digit of the binary expansion of s/p being 0 or 1, which allows us to determine the values of $\sigma(p, 2^j s)$ for ww' consecutive values of j in a single step by computing the corresponding ww' bits of the binary expansion of s/p . For this we define

```

246
247 1 modf b x p i = (b~i * x) mod p
248 2 digf b x p i = (b * modf b x p i) div p
249 3 dsgn d = if d ≠ 0 then -1 else 1
250 4 blbits i j = digf (2~(w * w')) s p i div 2~(j * w) mod 2~w
251

```

and rewrite the inner sum for multiples of ww' as

$$\sum_{0 \leq j < nww'} x^j \sigma(p, 2^j s) = \sum_{0 \leq j < nww'} x^j \text{dsgn}(\text{digf}(2, s, p, j))$$

We split the sum into chunks and blocks, and introduce the tabulation

$$\begin{aligned} &= \sum_{0 \leq k < w} x^k \sum_{0 \leq j < w'} x^{jw} \sum_{0 \leq i < n} x^{iww'} \text{dsgn}(\text{blbits}(i, j) ! k) \\ &= \sum_{0 \leq b < 2^w} \left(\left(\sum_{0 \leq k < w} x^k b_k \right) \cdot \left(\sum_{0 \leq j < w'} x^{jw} \text{tab}(j, b) \right) \right) \end{aligned}$$

where $\cdot ! k$ selects the k th bit of its argument, such that $\text{blbits}(i, j) ! k = \text{digf}(2, iww' + jw + k)$.

The table is defined as

$$\text{tab}(j, b) = \sum_{0 \leq i < n} \begin{cases} x^{iww'} & \text{if } b = \text{blbits}(i, j) \\ 0 & \text{otherwise} \end{cases}$$

Since n is typically much larger than w and 2^w , most of the computational work goes into computing the table, which our optimisations make cheap: in an outer loop, we obtain chunks of ww' bits of the binary expansion and incrementally maintain $x^{iww'}$. The inner loop only bit-shifts and masks these bits, indexes into the table, and adds $x^{iww'}$ to table

23:8 Verifying an Efficient Algorithm for Computing Bernoulli Numbers

```

1 harvey_tab_mod2 invp p p' x s n = doN {
2   let curs = s; curx ← to_mont1 32 p p' 1; xww' ← mont_exp 32 p p' x (w*w');
3   let tab = 0; (* Initialise table to all 0s *)
4   (curs,curx,tab) ← for 0 n (λi (curs,curx,tab). doN {
5     (cbits,curs) ← invp_push_div_mod1 invp curs p; (* Get 64 more bits of s/p *)
6     (_,tab) ← for 0 w' (λj (cbits,tab). doN {
7       let bbits = cbits AND (2^w-1); cbits = cbits>>w; (* Obtain block bits *)
8       (* Update table *)
9       v ← mop_tab_lookup tab j bbits; v ← mont_add_relaxed1 32 v curx;
10      tab ← mop_tab_upd tab j bbits v;
11      return (cbits,tab)
12    }) (cbits,tab);
13    curx ← mont_times1 32 p p' curx xww'; (* Maintain curx *)
14    return (curs,curx,tab)
15  }) (curs,curx,tab);
16  return tab }

```

■ **Figure 2** Computation of the inner-sum table

264 entries. Additionally, the implementation makes similar optimisations to those we described
265 in Sec. 3. In particular, we use the relaxed Montgomery form for the table entries to allow
266 for a mostly branch free implementation of the inner loop.⁵ Moreover, the size of the table
267 ($w'2^w$ double-width machine words, i.e. 16 KiB in our current setting) is chosen such that it
268 fits comfortably inside the L1d cache of a typical modern CPU.

269 To incrementally obtain the bits of the binary expansion, we maintain the value of
270 modf , and observe that $\text{modf}(b, x, p, i + 1) = (b \cdot \text{modf}(b, x, p, i)) \bmod p$. However, instead
271 of performing an expensive division by p , we apply a further optimisation: we define
272 $p^- = \lfloor 2^{128}/p \rfloor$. Then, for $p < 2^{32}$ and $x < p$, we have $\lfloor 2^{64}x/p \rfloor \in \{\lfloor p^-x/2^{64} \rfloor, \lfloor p^-x/2^{64} \rfloor + 1\}$.
273 Using this, we can reduce the division and modulo operation to a multiplication operation.
274 To be able to detect the +1 case, we limit $p < 2^{31}$, and define:

```

275
276 1 invp_push_div_mod invp x p =
277 2   let z = (invp*x)>>64; x' = (-z*p) AND (2^32-1) in
278 3   if x'<p then (z,x') else ((z+1) AND (2^64-1), (x'-p) AND (2^32-1))
279

```

280 Here, \gg denotes the left shift operation and AND the bitwise ‘and’. With this, we get:

```

281
282 1 lemma compute_next_bits:
283 2   assumes "invp = 2 ^ 128 div p" "curs = modf (2 ^ 64) s p i"
284 3   assumes "0 ≤ s" "s < p" "p<2^31"
285 4   shows "invp_push_div_mod 32 64 128 invp curs p
286 5     = (digf (2 ^ 64) s p i, modf (2 ^ 64) s p (i+1))"
287

```

288 This shows how to obtain the next digit and maintain the value of modf . Note that $\lfloor 2^{128}/p \rfloor$
289 is precomputed outside the loop. Figure 2 displays the algorithm for computing the table.
290 After some initialisation, the outer loop iterates over the chunks. In each iteration of the
291 outer loop, the binary expansion of s/p is maintained and another 64 bits of digits are

⁵ The computation of the binary expansion of s/p we describe below does use a branch to correct a possible ‘off-by-one’ error. However, this happens so rarely that due to branch prediction, the cost of this branch is negligible.

```

1 bernoulli_num_harvey_wrapper2 dfs p g op2 k = doN {
2   p' ← mk_mont_N1' p; (* Compute p' *)
3   r ← if p-1 dvd k then doN { (* Case 2: p-1 / k *)
4     t ← prod_list_exc2 p p' p dfs; mont_diff1 p (mont_zero1) t
5   } else doN { (* Case 1 *)
6     t ← if k < p - 1 then harvey_select2 p' p op2 k g
7     else doN { (* Apply range reduction *)
8       let k' = k mod (p - 1); t ← harvey_select2 p' p op2 k' g;
9       t2 ← to_mont1_relaxed p p' k; t ← mont_times1 p p' t t2;
10      kinv' ← to_mont1 p p' k'; t3 ← mont_inv1 p p' kinv'; mont_times1 p p' t t3
11    };
12    t' ← prod_list2 p p' dfs; mont_times1 p p' t t'; (* Multiply with D_k *)
13    r ← of_mont1 p p' r; return r } (* Convert to normal numbers *)

```

■ **Figure 3** Computation of a single remainder (slightly simplified)

obtained in `cbits`. The inner loop iterates over the blocks, obtains the current block bits, and updates the table. After the inner loop, the outer loop maintains the value `curx` of $x^{iww'}$.

3.3 Putting Things Together

Finally, we assemble the aligned inner sum, the computation of the remainder, and the fallback to the slower algorithm when $2^k \equiv 0$ (i.e. when $k \bmod \text{ord}_p(2) = 0$) to obtain the algorithm `harvey_select2` with the following correctness theorem:

```

298
299 1 lemma harvey_select2_correct:
300 2   assumes "mont_ctxt_refine' W p p' ctxt"
301 3     and "even k" "¬(p-1) dvd k" "k ∈ {2..p - 1}" "prime p" "5 ≤ p"
302 4     and "op2 = ord p 2" "g < p" "residue_primroot p g"
303 5   shows "harvey_select2 p' p op2 k g
304 6     ≤ ↓(mont_rel ctxt) (SPEC (λr. [r = bernoulli_rat k] (qmod p)))"
305

```

That is, under the listed preconditions, `harvey_select2` will return the Montgomery form of an integer r with $r \equiv_p B_k$.

Next, we address some of the preconditions. First, we compute $p' = p^{-1} \bmod 2^W$ via Hensel lifting. Next, as mentioned before, when $(p-1) \mid k$ we have $N_k \equiv_p -D_k/p$, which allows us to easily determine $N_k \bmod p$ directly. Moreover, when $k \geq p-1$, we apply the previously-mentioned range reduction based on Kummer's congruence. Finally, we multiply the result with D_k to obtain $N_k \bmod p$, and convert the result from Montgomery form to normal numbers.

The resulting algorithm is displayed in Fig. 3. In addition to the parameters p , g , $op2$, and k , it also takes the prime factorisation `dfs` of D_k . The function `prod_list_exc2` computes D_k/p by multiplying all prime factors except p . The correctness theorem is:

```

317
318 1 lemma bernoulli_num_harvey_wrapper2_correct:
319 2   assumes "even k" "k ≠ 0" "prime p" "3 ≤ p" "p < 2^31"
320 3   assumes "dfs = denom_factors k" "residue_primroot p g" "g < p" "op2 = ord p 2"
321 4   shows "bernoulli_num_harvey_wrapper2 dfs p g op2 k
322 5     ≤ SPEC (λr. r = bernoulli_num k mod p)"
323

```

In words: the algorithm returns $N_k \bmod p$ for any even natural number $k \neq 0$ and any odd prime less than 2^{31} .

23:10 Verifying an Efficient Algorithm for Computing Bernoulli Numbers

```

1 compute_primes_mods_est t d k = doN {
2   (ok,Y) ← estimate_check_a_priori_bound k; (* A-priori bound *)
3   if ¬ok then return (False, (0, [], [])) (* Out-of-bounds *)
4   else doN {
5     (fc,pc) ← mk_factor_prime_cache (max (k+2) Y); (* Prime sieve *)
6     dfs ← filter_denom_factors pc k; (* Get factorisation of D_k *)
7     (_,pc) ← adjust_primes2 dfs k pc; (* Precise bound *)
8     pgos ← compute_harvey_pgos pc fc; (* Compute generator and ord 2 p *)
9     result ← pm.pmap_par_array (max 2 t) d (k,dfs) pgos (length pgos);
10    return (True, (length result,result,dfs)) } } (* Ret result and D_k factors *)

```

■ **Figure 4** Parallel computation of the prime numbers and remainders

326 **4** Computing B_k as a Rational Number

327 The full algorithm consists of a preprocessing phase to create a large array of independent
328 ‘tasks’, each being represented with a triple $(p, g_p, \text{ord}_p(2))$, then the main part of computing
329 $B_k \bmod p$ for each of these, and finally a modular reconstruction phase to determine B_k as a
330 rational number. Figure 4 shows a high-level view of the algorithm without the modular
331 reconstruction phase.

332 First, we compute a rough a-priori upper bound Y for the largest prime that will be
333 needed to gather enough modular information to reconstruct B_k . We also check that k is
334 small enough to not cause overflows later, and otherwise return an error. Next, we use a
335 simple sieving algorithm to compute all primes up to Y and a factor map fc that maps
336 every composite number to a prime factor. Next, we compute the prime factors of D_k using
337 the von Staudt–Clausen theorem. We then use this information to determine a good upper
338 bound for $\log_2 |N_k|$ and drop the unnecessary primes from the list. For each remaining prime
339 p , we then compute $\text{ord}_p(2)$ and find a generator g_p of $\mathbb{Z}/p\mathbb{Z}$. We refer to the list of triples
340 $(p, g_p, \text{ord}_p(2))$ as the ‘task list’ (or, in the Isabelle formalisation, the ‘pgo list’). Finally, we
341 compute the modular information for each of the primes in parallel and return the result as
342 well as the prime factorisation of D_k .

343 In the following sections, we will explain each of these steps as well as the modular
344 reconstruction step in some detail.

345 **4.1 A-priori bound**

346 The function `estimate_check_a_priori_bound k` (cf. Fig. 4) returns a pair $(ok, Y + 1)$. If
347 $ok = \text{True}$, then Y is an upper bound for the largest prime number that will be needed
348 to reconstruct N_k . If $ok = \text{False}$, we cannot guarantee that our 32-bit arithmetic will not
349 overflow, and we abort the algorithm. The first ingredient to compute Y is a bound on B_k :

350 ► **Theorem 8** (Upper bound for B_k). *Let $k \geq 4$ be an even integer. Then*

$$351 \quad \log_2 |B_k| \leq (k + \frac{1}{2}) \log_2 k - c_1 k + c_2$$

352 *where $c_1 = \frac{1}{\ln 2} + \log_2 \pi + 1 \approx 4.0942$ and $c_2 = \frac{3}{2} + \frac{9}{2} \log_2 \pi - \log_2 90 + \frac{1}{48 \ln 2} \approx 2.4699$.*

353 **Proof.** We use Theorem 2 to express $|B_k|$ in terms of $\zeta(k)$. The monotonicity of $\zeta(x)$ for
354 real $x > 1$ implies $\zeta(k) \leq \zeta(4) = \pi^4/90 \approx 1.082$. Combining this with Stirling’s inequality
355 $\ln(k!) \leq \frac{1}{2} \ln(2\pi k) + k \ln k - k + \frac{1}{12k}$ and noting that $\frac{1}{12k} \leq \frac{1}{48}$ yields the result. ◀

356 We can convert this into a bound on N_k by combining it with the von Staudt–Clausen
 357 theorem: it implies $D_k \mid 2(2^k - 1)$ and therefore $\log_2 D_k < k + 1$. Lastly, a weak (but
 358 non-asymptotic) version of the Prime Number Theorem states that $\sum_{p \leq x} \ln p \geq 0.82x$ for
 359 all $x \geq 97$ [6]. This was also formalised specifically for this project. Combining all of these
 360 bounds, we obtain the following:

361 ► **Theorem 9** (A-priori prime bound). *Let $k > 0$ be an integer. Then $|N_k| < \prod_{p \leq Y} p$ where*

$$362 \quad Y = \max(97, c_1(k + \frac{1}{2}) - c_2k + c_3)$$

363 *for constants $c_1 = \frac{625}{738} \approx 0.8469$, $c_2 = \frac{50941792385100000}{19440408116330496} \approx 2.6204$, $c_3 = \frac{23688125}{8060928} \approx 2.9386$.*

364 In our algorithm, we use the following slightly weaker upper bound:

$$365 \quad Y = \max\left(97, \left\lceil \frac{217k \lceil \log_2 k \rceil}{256} \right\rceil + \left\lceil \frac{217 \lceil \log_2 k \rceil}{512} \right\rceil + 3 - \left\lfloor \frac{335k}{128} \right\rfloor\right)$$

366 If $Y \geq 2^{31}$, we return $ok = \text{False}$. We prove that Y will be in bounds for $k \leq 105,946,388$.
 367 We leave it to future work to delay the check until after we have pruned the list of primes
 368 with the more precise bound.

369 Note that our bound Y is lower than Harvey’s since, first, we can make use of all
 370 primes $\leq Y$ because we compute N_k instead of B_k , and secondly we use a somewhat sharper
 371 inequality for $\prod_{p \leq x} p$.

372 4.2 Sieving

373 The prime sieving algorithm `mk_factor_prime_cache` K returns a tuple (fc, pc) , where pc
 374 (“prime cache”) is the list of primes p with $2 < p < K$, and fc (“factor cache”) encodes
 375 a mapping that maps any integer between 0 and K to its smallest prime factor (if it is
 376 composite) or to 0 (if it is prime or ≤ 1). The sieving algorithm starts by mapping all
 377 numbers to 0 and then proceeds in the usual fashion. When sieving is finished, the list of
 378 primes is extracted from the factor cache.

379 The advantage of the factor cache is that it allows us to not only quickly list prime
 380 numbers and determine whether a number is prime, but also to factor numbers efficiently.
 381 This speeds up the `compute_harvey_pgos` algorithm (cf. Sec 4.4).

382 4.3 Computing the Denominator and Precise Bounding

383 We can now determine the prime factorisation of D_k by simply taking all primes $p \leq k$ with
 384 $(p - 1) \mid k$. The algorithm `adjust_primes2 dfs k pc` sums up $\log_2 p$ for all of these to get an
 385 approximation of $\log_2 D_k$ that easily fits in a machine integer. Together with our estimate
 386 for $\log_2 |B_k|$, we can now compute a relatively precise estimate of $\log_2 |N_k|$.

387 We then choose a prefix P of the sieved primes such that $\sum_{p \in P} \log_2(p) \geq \log_2 |N_k|$. For
 388 this, we use a fixed-point approximation of $\log_2 p$. This avoids using floating-point numbers
 389 or arbitrary precision integers and is still precise enough in practice. Note that we do not
 390 actually store the prime 2 in the sieve or pruned list of primes, but still consider it for the
 391 bound. Before reconstructing $|N_k|$ (cf. Sec. 4.6), we do add the congruence $N_k \equiv_2 1$ (which
 392 holds for any even k) to our modular information.

393 4.4 Preparing the Task List

394 The algorithm `compute_harvey_pgos pc fc` takes the pruned prime list and the factor map
 395 and computes a list of task entries of the form $(p, g, o, 0)$. Here, p is the prime itself, g is a

23:12 Verifying an Efficient Algorithm for Computing Bernoulli Numbers

```

1 lemma pm.pmap_par_array_correct:
2   assumes "n=length pgos" "1<t" "∀pgo∈set pgos. bernpre (k,dfs) pgo"
3   shows "pm.pmap_par_array t d (k,dfs) pgos n
4     ≤ SPEC (λpgos'. list_all2 (bernspec (k,dfs)) pgos pgos')"

```

where the pre- and postconditions are defined as:

```

1 bernpre (k,dfs) (p,g,op2,rX) =
2   prime p ∧ p∈{3..2^(num_len-1)} ∧ residue_primroot p g
3   ∧ g < p ∧ op2 = ord p 2 ∧ dfs = denom_factors k
4 bernspec (k,dfs) (p,g,op2,rX) (p',g',op2',r) =
5   (p',g',op2') = (p,g,op2) ∧ r = bernoulli_num k mod (int p)

```

■ **Figure 5** Correctness lemma for parallel computation of the modular information. The input list contains tuples of the form $(p, g, op2, rX)$, where the last element rX is the not yet specified result. In the result list, the last element is replaced by the correct result, and the other elements are unchanged. Here, `list_all2` is the natural relator on lists, relating lists of the same length element-wise. Note that the parameters t and d are used to control the level of parallelisation and do not affect the result.

396 generator of $\mathbb{Z}/p\mathbb{Z}$, o is the multiplicative order of 2 in $\mathbb{Z}/p\mathbb{Z}$, and 0 is a placeholder that will
397 be filled in by Harvey's algorithm.

398 A number g is a generator of $\mathbb{Z}/p\mathbb{Z}$ iff it has maximal order, i.e. if $\text{ord}_p(g) = p - 1$. We
399 can determine whether g is a generator by checking that $g^{(p-1)/q} \not\equiv_p 1$ for all prime factors q
400 of $p - 1$. We find the smallest generator g_p by simply checking $g = 2, 3, \dots, p - 1$ and taking
401 the first one that works.⁶

402 As for computing the order: $\text{ord}_p(2)$ is defined as the smallest positive integer n such
403 that $2^n \equiv_p 1$. Lagrange's theorem tells us that $\text{ord}_p(2) \mid p - 1$. Therefore, every factor q
404 in the prime factorisation of $\text{ord}_p(2)$ must also be a factor of $p - 1$, and we can determine
405 its multiplicity $\nu_q(\text{ord}_p(2))$ by letting $i \leftarrow 0$, $x \leftarrow 2^{(p-1)/q^e} \bmod p$ (where $e = \nu_q(p - 1)$)
406 and then repeating $i \leftarrow i + 1$, $x \leftarrow x^q \bmod p$ until $x = 1$, at which point $i = \nu_q(\text{ord}_p(2))$ as
407 desired.

408 We are not aware of better algorithms for computing the order and finding a generator;
409 however, even the worst-case running time of these is negligible compared to the main part
410 of our algorithm.

411 4.5 Parallel Map

412 To compute the modular information in parallel, we instantiate a generic parallel map
413 combinator with `bernoulli_num_harvey_wrapper2` (cf. Sec. 3.3). We elide the boilerplate code
414 for the instantiation, and only display the resulting correctness lemma in Fig. 5.

415 4.6 Postprocessing

416 Figure 6 displays our Isabelle formalisation of Harvey's full algorithm. After handling some
417 special cases for $k = 0$, $k = 1$, and odd k , it computes the modular information and a prime

⁶ Empirically, g_p is almost always quite small. For the first 10^6 primes, the mean and maximum of g_p are 4.91 and 94, with $g_p \in \{2, 3\}$ in 60% of the cases. It is known that $g_p \in O(p^{\frac{1}{4}+\epsilon})$ and the generalised Riemann hypothesis implies $g_p \in O(\log^6 p)$.

```

1 "bern_crt t d k = doN {
2   if k=0 then return (True,1,1) else if k=1 then return (True,-1,2)
3   else if odd k then return (True,0,1)
4   else doN { (* Regular case: even k, k>1 *)
5     (ok, (n,pgors,dfs)) ← compute_primes_mods_est t d k; (* Compute mod. info *)
6     if ¬ok then return (False,0,0) (* Out-of-bounds *)
7     else doN {
8       (n,ams) ← map_to_ams2 n pgors; (* Map to (r,p)-pairs, add (1,2) *)
9       (num,m) ← crt n ams; (* Chinese remaindering *)
10      let num = (if 4 dvd k then num-m else num); (* Sign adjustment *)
11      let denom = (∏x←dfs. x); (* Compute D_k *)
12      return (True,num,denom)
13    } } }"

```

■ **Figure 6** The abstract version of Harvey’s algorithm in Isabelle

418 factorisation of the denominator. If the *ok* flag is false, we abort the algorithm. Otherwise,
 419 the (p, g, o, r) tuples in *pgors* are converted to (r, p) pairs in *ams*. We also add the pair $(1, 2)$
 420 since N_k is always odd. At this point, we know that for each pair $(r, p) \in \text{ams}$, we have
 421 $N_k \bmod p = r$. Moreover, the primes P in these pairs are disjoint, and their product is
 422 greater than N_k . Chinese remaindering gives us a pair (num, m) with $\text{num} = N_k \bmod m$
 423 and $m = \prod_{p \in P} p$. Thus, $N_k = \text{num}$ if $N_k \geq 0$ and $N_k = \text{num} - m$ if $N_k < 0$, and we know
 424 that $N_k < 0$ iff $4 \mid k$. Finally, we compute the denominator $\text{denom} = D_k$ from its prime
 425 factorisation.

426 4.6.1 Fast Chinese Remaindering

427 The Chinese Remainder Theorem (CRT) allows us to determine $N_k \bmod (\prod_{p \in P} p)$ from the
 428 values of $N_k \bmod p$ for every $p \in P$.

429 A naïve implementation of the CRT is too slow for numbers as big as ours. We therefore
 430 follow the approach called *remainder trees* outlined in the classic textbook by von zur Gathen
 431 and Gerhard [23]. The basic data structure is a binary tree where every leaf has a pair of
 432 modular data and modulus of the form $(y \bmod p, p)$ attached to it, and every internal node
 433 has a modulus attached to it that is the product of the moduli of its children. Given this
 434 tree t , we can compute our desired result $y \bmod (\prod_{p \in P} p)$ as $f(\prod_{p \in P} p, t)$, where f is the
 435 following simple recursive function:

$$\begin{aligned}
 436 \quad f(M, \text{Leaf}(x, p)) &= (x \cdot ((M \bmod m^2)/p)^{-1}) \bmod p \\
 437 \quad f(M, \text{Node}(l, m, r)) &= f(M \bmod m^2, l) \cdot \text{root}(r) + f(M \bmod m^2, r) \cdot \text{root}(l)
 \end{aligned}$$

438 This can also be parallelised easily since different branches of the tree are independent.

439 While the abstract version of our algorithm uses a tree data structure, which is initially
 440 created from the list of input pairs, we later refine this to use the list of input pairs directly
 441 for the leafs, and another array for the inner nodes. This implementation is slightly more
 442 memory-efficient than using a data structure with explicit pointers. However, in practice, we
 443 expect memory usage to be dominated by the space for the integers.

444 4.6.2 Arbitrary Precision Integers

445 While all other parts of our algorithm use machine words, CRT and the computation of D_k
 446 require large integers. We take a pragmatic approach and import the GNU Multiprecision

23:14 Verifying an Efficient Algorithm for Computing Bernoulli Numbers

447 Library (GMP) [11] into our formalisation. We declare an assertion `mpzA` for big integers,
448 and the GMP library functions and use Isabelle’s specification mechanism to specify the
449 Hoare-triples for them. We then instruct Isabelle LLVM to translate calls to the specified
450 functions to the corresponding symbols from the GMP library.

451 The specification mechanism ensures that we do not know more about the constants than
452 what was specified. Any such knowledge is dangerous, as it could be used to prove behaviour
453 that is not exhibited by the GMP library. At the same time, the specification mechanism
454 requires us to prove that a model for the specification exists. This is an important sanity
455 check, as it prevents us from accidentally specifying contradictory statements.

456 **5 Code Export and Final Correctness Theorem**

457 Finally, we use Sepref to refine `bern_crt` (cf. 4.6) to Isabelle LLVM and prove:

```
458  
459 1 (bern_crt_impl, bern_crt) : sn64A * sn64A * u32A → u1A × mpzA × mpzA  
460
```

461 Note that Isabelle LLVM uses non-negative signed integers in various places, which are
462 related to natural numbers by the `sn64A` refinement assertion. To make the algorithm usable
463 from C++, we define the wrapper function `bern_crt_impl_wrapper` that returns void and
464 uses pointer parameters to pass the results instead. Combining the refinement lemmas yields
465 the following correctness theorem:

```
466  
467 1 theorem bern_crt_impl_wrapper_correct: "llvm_htriple (  
468 2   ll_pto okX okp * ll_pto numX nump * ll_pto denomX denomp  
469 3   * sn64A t ti * sn64A d di * u32A k ki)  
470 4   (bern_crt_impl_wrapper okp nump denomp ti di ki)  
471 5   (λ_. EXS ok numi denomi num denom.  
472 6     ll_pto ok okp * ll_pto numi nump * ll_pto denomi denomp  
473 7     * mpzA num numi * mpzA denom denomi  
474 8     * (k ≤ 105946388 → ok ≠ 0) * (ok ≠ 0 → num = N k ∧ denom = D k)
```

476 The preconditions of this Hoare-triple are that `okp`, `nump`, and `denomp` point to valid memory,
477 that `ti` and `di` are 64-bit non-negative signed integers, and that `ki` is a 32-bit unsigned with
478 value `k`. Then, after calling `bern_crt_impl_wrapper okp nump denomp ti di ki`, the pointers
479 `okp`, `nump`, and `denomp` will point to the values `ok`, `numi`, and `denomi`; `numi` and `denomi` will be
480 valid GMP numbers representing the integers `num`, and `denom`; and if $k \leq 105,946,388$, `ok` will
481 be non-zero, and if `ok` is non-zero, `num` and `denom` represent the Bernoulli number B_k . Note
482 that we explicitly specified N_k and D_k here, as this also guarantees that `num` and `denom` have
483 no common divisors.

484 This correctness theorem only depends on the Isabelle LLVM model, the formalisation of
485 separation logic, our specification of GMP, and the definition of Bernoulli numbers. All the
486 intermediate refinement steps need not to be trusted.

487 Finally, we use Isabelle LLVM to export the function to actual LLVM text and to generate
488 a header file to interface with the generated code from C/C++.

```
489  
490 1 export_llvm bern_crt_impl_wrapper is "void bern_crt (  
491 2   uint8_t *, gmp_mpz_struct**, gmp_mpz_struct**, sint64_t, sint64_t, uint32_t)"  
492 3   file "../llvm_export/bern_crt.ll"
```

494 We then implement a command line interface in C++, with which we run our experiments
495 to generate actual Bernoulli numbers. For the more detailed timing measurements, we
496 instrumented the LLVM code to invoke callbacks into our C++ program when certain phases
497 of the algorithm are completed.

Input k	10^5	$\lceil 10^{5.5} \rceil$	10^6	$\lceil 10^{6.5} \rceil$	10^7	$\lfloor 10^{7.5} \rfloor$	10^8
time elapsed (s)	1.12	1.41	5.75	44.6	432.3	4,215	45,458
CPU time (s)	12.56	61.45	430.11	4,318.58	46,960	475,160	5,330,806
CPU factor	12.1	44.0	74.9	96.9	108.6	112.7	117.3
memory (MB)	48.9	124.4	378.8	1220	3,948	11,775	36,769
sieving	0.4%	1.3%	1.1%	0.7%	0.3%	0.1%	0.0%
generator/ord(2)	6.5%	18.9%	15.6%	7.3%	2.7%	0.9%	0.3%
main phase	25.8%	61.5%	64.6%	81.2%	92.4%	96.9%	98.9%
reconstruction	30.5%	8.4%	8.2%	4.5%	1.9%	0.7%	0.3%
writing result	2.1%	8.7%	10.0%	6.2%	2.7%	1.3%	0.5%

■ **Figure 7** Upper half: Performance statistics of the exported LLVM code as given by the GNU `time` command. The CPU factor is an indicator for the amount of actual parallelism (best possible would be 128); Lower half: Percentage of elapsed time spent in each part of the algorithm. Note that the “writing” phase is conducted by an (unverified) GMP routine and contains both the conversion of the GMP integer into a base-10 string and the actual writing of that string to the result file.

6 Benchmarks

Table 7 shows the performance data of the exported LLVM code, compiled with Clang 18.1.8 and running on a single “standard” node of the LEO5 cluster operated by the University of Innsbruck with two 2.60 GHz Intel Xeon Platinum 8358 CPUs with 32 cores each and hyperthreading.

Among other things, the table shows the percentage of elapsed time spent in the various parts of the algorithm. It can be seen that especially for larger inputs, the main phase (i.e. computing the modular information $B_k \bmod p$ for each prime p) dominates the running time so that optimisation of the other parts will not result in a noticeable speed-up.

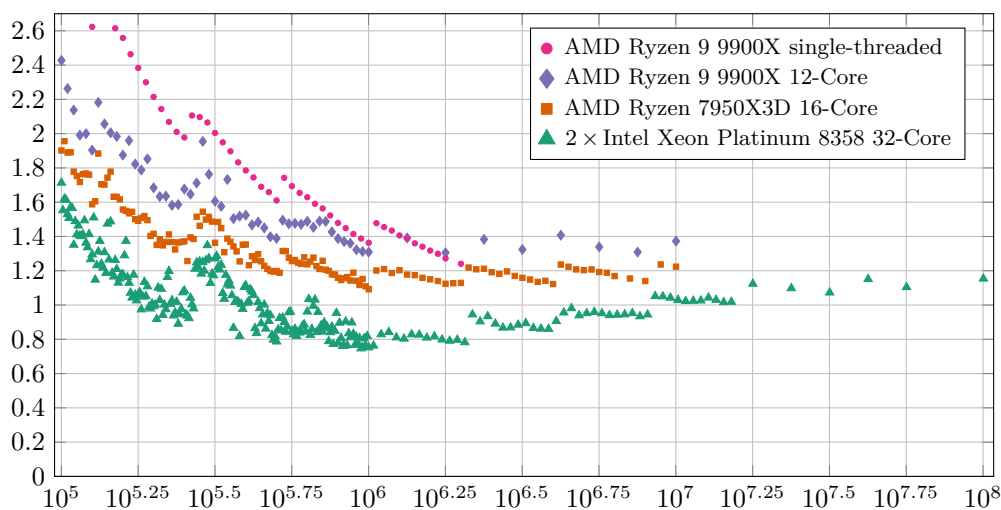
Figure 8 shows the relative performance of our verified LLVM implementation and Harvey’s unverified C++ implementation on two different machines: a 12-core desktop computer, a 16-core server, and a cluster node with two 32-core server CPUs. For the desktop computer, we also show the single-core performance.

The data suggests that our implementation is significantly slower than Harvey’s for small inputs and competitive with it on the more interesting larger inputs (roughly $k \geq 3 \cdot 10^5$). Interestingly, our implementation slightly but consistently outperforms Harvey’s on the cluster in the range $3 \cdot 10^5 \leq k \leq 5 \cdot 10^6$. This may be due to the different implementation of parallelisation.

Finding out precisely where our implementation is slower than Harvey’s and why will require thorough microbenchmarking of various parts of the code.

7 Related Work

We are not aware of any other work on the verification of algorithms involving Bernoulli numbers. As far as proof assistants are concerned, we are only aware of two others that even have a *definition* of Bernoulli numbers, namely HOL Light and Lean. However, neither system’s library seems to have an algorithm for computing them other than using the recurrence that follows directly from their definition. We are not aware of any verification efforts for Bernoulli number algorithms outside proof assistants either.



■ **Figure 8** The ratio of the elapsed time of our verified implementation divided by that of Harvey's, for a range of input indices k and on several different machines.

525 As for unverified algorithms, there are various approaches apart from Harvey's 2010
 526 algorithm (which we formalised in this paper) that are more efficient than the naïve approach:
 527 ■ as the Akiyama–Tanigawa transform of the sequence $(\frac{1}{k+1})_{k \in \mathbb{N}}$ [2, 14]
 528 ■ by computing D_k and approximating $\zeta(k)$ to sufficient precision
 529 ■ by computing the closely-related *tangent numbers*, e.g. by performing the division
 530 $\sin(z)/\cos(z)$ as formal power series, either directly or via Kronecker substitution (cf. e.g.
 531 Brent and Harvey [3])


532 Harvey discusses some of these in some more detail in his paper. The first of them is already
 533 available in Isabelle [4], but it is quite inefficient for even moderate values of k . The last two
 534 have the same asymptotic running time as the algorithm we verify, namely $O(k^2 \log^{2+o(1)} k)$,
 535 but for various reasons discussed in Harvey's paper, they seem to perform much worse
 536 in practice. Note however that the third algorithm computes not only B_k but the entire
 537 sequence B_0, \dots, B_k .

538 In 2012, Harvey also published another multi-modular algorithm for computing B_k which
 539 achieves subquadratic running time by computing B_k modulo prime powers instead of simply
 540 primes. Asymptotically, this is a large improvement over all previously known algorithms;
 541 however, to date, there is still no implementation, and it is not clear how well it would
 542 perform for feasible input sizes in practice. In fact, it seems that the only fast algorithms
 543 that have actual implementations are the ones based on approximating $\zeta(k)$ and Harvey's
 544 2010 algorithm.

545 The Isabelle LLVM framework has been used for verifying efficient implementations of
 546 several tools and algorithms, mainly for SAT solving [10, 18] and model checking [16, 12].
 547 Its parallel extension has only been used for sorting algorithms [19] so far.

548 8 Conclusion

549 We fully verified a challenging state-of-the-art mathematical algorithm to compute Bernoulli
 550 numbers, achieving performance comparable to highly optimised hand-written C++ code. A
 551 large amount of machinery had to be developed to achieve this, including:



■ **Figure 9** Single-core performance of Harvey's implementation and our implementation.

552 **Mathematical background:** bounds for Bernoulli numbers, the Voronoi and Kummer congru-
 553 ences, non-asymptotic bounds for the Chebyshev ϑ function ($\vartheta(x) = \sum_{p \leq x} \ln p \geq 0.82x$
 554 for all $x \geq 97$)

555 **Algorithms, abstractly and verified down to LLVM:** prime sieving, Chinese remaindering
 556 via remainder trees, computing $\lceil \log_2 n \rceil$, fixed-point approximations of $\log_2 n$, efficient
 557 computation of the binary fraction expansion of a rational number, Hensel lifting, the
 558 extended Euclidean algorithm, Montgomery ('REDC') arithmetic, computing $\text{ord}_{\mathbb{Z}/n\mathbb{Z}}(x)$,
 559 finding generators in $\mathbb{Z}/n\mathbb{Z}$

560 **Extending the Isabelle/LLVM framework:** axiomatisation of GMP integers, destructive par-
 561 allel maps over an array, low-level bit arithmetic

562 We expect many of these developments to be useful for other applications.

563 The total development is about 30 kLOC in Isabelle. We also wrote prototype algorithms
 564 in C++ to understand the algorithm and its critical optimisations.

565 ——— References ———

- 566 1 J. C. Adams. On the calculation of Bernoulli's numbers up to B_{62} by means of Staudt's theorem.
 567 In *Report of the meeting of the British Association for the Advancement of Science*, Rep. Brit.
 568 Assn., pages 8–15, 1877. URL: <https://gallica.bnf.fr/ark:/12148/bpt6k78160g>.
- 569 2 Shigeki Akiyama and Yoshio Tanigawa. Multiple zeta values at non-positive integers. *The*
 570 *Ramanujan Journal*, 5(4):327–351, 2001. doi:10.1023/A:1013981102941.
- 571 3 Richard P. Brent and David Harvey. *Fast Computation of Bernoulli, Tangent and Secant*
 572 *Numbers*, pages 127–142. Springer New York, 2013. doi:10.1007/978-1-4614-7621-4_8.
- 573 4 Lukas Bulwahn and Manuel Eberl. Bernoulli numbers. *Archive of Formal Proofs*, January
 574 2017. <https://isa-afp.org/entries/Bernoulli.html>, Formal proof development.
- 575 5 Henri Cohen. *Number Theory: Volume II: Analytic and Modern Tools*. Graduate Texts in
 576 Mathematics. Springer New York, 2007.
- 577 6 Manuel Eberl. Concrete bounds for Chebyshev's prime counting functions. *Archive of*
 578 *Formal Proofs*, September 2024. https://isa-afp.org/entries/Chebyshev_Prime_Bounds.html, Formal proof development.
- 579 7 Manuel Eberl. Kummer's congruence. *Archive of Formal Proofs*, March 2024. https://isa-afp.org/entries/Kummer_Congruence.html, Formal proof development.

- 582 8 Eugene Eric Kim and Betty Alexandra Toole. Ada and the first computer. *Scientific American*,
583 280(5):76–81, May 1999. doi:10.1038/scientificamerican0599-76.
- 584 9 Greg Fee and Simon Plouffe. An efficient algorithm for the computation of Bernoulli numbers,
585 2007. URL: <https://arxiv.org/abs/math/0702300>, arXiv:math/0702300.
- 586 10 Mathias Fleury and Peter Lammich. A more pragmatic CDCL for isasat and targetting LLVM
587 (short paper). In Brigitte Pientka and Cesare Tinelli, editors, *Automated Deduction - CADE*
588 *29 - 29th International Conference on Automated Deduction, Rome, Italy, July 1-4, 2023,*
589 *Proceedings*, volume 14132 of *Lecture Notes in Computer Science*, pages 207–219. Springer,
590 2023. doi:10.1007/978-3-031-38499-8_12.
- 591 11 The GNU multiple precision arithmetic library. URL: <https://gmpilib.org/>.
- 592 12 Arnd Hartmanns, Bram Kohlen, and Peter Lammich. Efficient formally verified maximal
593 end component decomposition for mdps. In André Platzer, Kristin Yvonne Rozier, Matteo
594 Pradella, and Matteo Rossi, editors, *Formal Methods - 26th International Symposium, FM*
595 *2024, Milan, Italy, September 9-13, 2024, Proceedings, Part I*, volume 14933 of *Lecture Notes*
596 *in Computer Science*, pages 206–225. Springer, 2024. doi:10.1007/978-3-031-71162-6_11.
- 597 13 David Harvey. A multimodular algorithm for computing Bernoulli numbers. *Math. Comput.*,
598 79(272):2361–2370, 2010. doi:10.1090/S0025-5718-2010-02367-1.
- 599 14 M. Kaneko. The Akiyama–Tanigawa algorithm for Bernoulli numbers. *Journal of Integer*
600 *Sequences*, 3, 2000.
- 601 15 Donald E. Knuth and Thomas J. Buckholtz. Computation of tangent, Euler, and
602 Bernoulli numbers. *Mathematics of Computation*, 21(100):663–688, 1967. doi:10.1090/
603 s0025-5718-1967-0221735-9.
- 604 16 Bram Kohlen, Maximilian Schäffeler, Mohammad Abdulaziz, Arnd Hartmanns, and Peter
605 Lammich. A formally verified IEEE 754 floating-point implementation of interval iteration for
606 mdps. *CoRR*, abs/2501.10127, 2025. URL: <https://doi.org/10.48550/arXiv.2501.10127>,
607 arXiv:2501.10127, doi:10.48550/ARXIV.2501.10127.
- 608 17 Peter Lammich. Generating verified LLVM from isabelle/hol. In John Harrison, John O’Leary,
609 and Andrew Tolmach, editors, *10th International Conference on Interactive Theorem Proving,*
610 *ITP 2019, September 9-12, 2019, Portland, OR, USA*, volume 141 of *LIPICs*, pages 22:1–22:19.
611 Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.ITP.2019.22.
- 612 18 Peter Lammich. Fast and verified UNSAT certificate checking. In Christoph Benzmüller,
613 Marijn J. H. Heule, and Renate A. Schmidt, editors, *Automated Reasoning - 12th International*
614 *Joint Conference, IJCAR 2024, Nancy, France, July 3-6, 2024, Proceedings, Part I*, volume
615 14739 of *Lecture Notes in Computer Science*, pages 439–457. Springer, 2024. doi:10.1007/
616 978-3-031-63498-7_26.
- 617 19 Peter Lammich. Refinement of parallel algorithms down to LLVM: applied to practically efficient
618 parallel sorting. *J. Autom. Reason.*, 68(3):14, 2024. doi:10.1007/S10817-024-09701-W.
- 619 20 Peter Lammich and Thomas Tuerk. Applying data refinement for monadic programs to
620 Hopcroft’s algorithm. In *Proc. of ITP*, volume 7406 of *LNCS*, pages 166–182. Springer, 2012.
- 621 21 D. H. Lehmer. An extension of the table of Bernoulli numbers. *Duke Mathematical Journal*,
622 2(3), September 1936. doi:10.1215/s0012-7094-36-00238-7.
- 623 22 Peter L. Montgomery. Modular multiplication without trial division. *Mathematics of Compu-*
624 *tation*, 44(170):519–521, 1985. doi:10.1090/s0025-5718-1985-0777282-x.
- 625 23 Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. Cambridge
626 University Press, 3rd edition, 2013.
- 627 24 Alexander Yee and Shigeru Kondo. 10 trillion digits of pi: A case study of summing hyper-
628 geometric series to high precision on multicore systems. Technical report, Siebel School of
629 Computing and Data Science, 2011. URL: <https://hdl.handle.net/2142/28348>.