

A MULTIMODULAR ALGORITHM FOR COMPUTING BERNOULLI NUMBERS

DAVID HARVEY

ABSTRACT. We describe an algorithm for computing Bernoulli numbers. Using a parallel implementation, we have computed B_k for $k = 10^8$, a new record. Our method is to compute B_k modulo p for many small primes p and then reconstruct B_k via the Chinese Remainder Theorem. The asymptotic time complexity is $O(k^2 \log^{2+\varepsilon} k)$, matching that of existing algorithms that exploit the relationship between B_k and the Riemann zeta function. Our implementation is significantly faster than several existing implementations of the zeta-function method.

1. INTRODUCTION

The Bernoulli numbers B_0, B_1, B_2, \dots are rational numbers defined by

$$\frac{x}{e^x - 1} = \sum_{k=0}^{\infty} \frac{B_k}{k!} x^k.$$

In this paper we focus on the problem of computing B_k , as an exact rational number, for a single large (even) value of k . To date, the most efficient algorithms for computing an isolated B_k have been based on the formula

$$(1) \quad B_k = (-1)^{k/2+1} \frac{2\zeta(k)k!}{(2\pi)^k}, \quad k \geq 2 \text{ and } k \text{ even,}$$

where $\zeta(s)$ is the Riemann zeta function [IR90, p. 231]. One first computes a high-precision approximation for B_k via (1), using the Euler product for $\zeta(s)$. The exact denominator of B_k is known by the von Staudt–Clausen theorem. Provided that the numerical approximation to B_k has enough correct bits — it turns out that $O(k \log k)$ bits are enough — we may recover the numerator of B_k exactly. We will refer to this algorithm and its variants as the ‘zeta-function algorithm’.

The zeta-function algorithm has been rediscovered numerous times. The formula (1) itself goes back to Euler. Chowla and Hartung [CH72] give an ‘exact formula’ for B_k , the evaluation of which is tantamount to executing the algorithm, but using the defining sum for $\zeta(k)$ instead of the Euler product and using $2(2^k - 1)$ in place of the exact denominator of B_k . Fillebrown [Fil92] mentions [CH72] and points out that using the Euler product is asymptotically faster; she gives a detailed time and space complexity analysis and attributes the algorithm to Herbert Wilf (unpublished). Fee and Plouffe [FP07] state that they discovered and implemented the zeta-function algorithm in 1996. Kellner [Kel02] also refers to [CH72] and suggests

Received by the editor November 17, 2008.

2010 *Mathematics Subject Classification*. Primary 11B68, Secondary 11Y60.

©2010 American Mathematical Society
Reverts to public domain 28 years from publication

using the Euler product. He mentions some data produced by Fee and Plouffe, although not their algorithm. According to Pavlyk [Pav08], the implementation in Mathematica 6 [Wol07] derives from Kellner's work. Stein [Ste07, p. 30] indicates another independent discovery by Henri Cohen, Karim Belabas, and Bill Daly (unpublished) that led to the implementation in PARI/GP [Par07].

In this paper we present a new algorithm for computing B_k that does not depend on (1) or any properties of $\zeta(s)$. The basic idea is to compute B_k modulo p for sufficiently many small primes p and then reconstruct B_k using the Chinese Remainder Theorem. Using a parallel implementation, we have computed B_k for $k = 10^8$, improving substantially on the previous record of $k = 10^7$ [Pav08].

Let $M(n)$ denote the time (bit operations) required to multiply two n -bit integers. We may assume that $M(n) = O(n \log^{1+\varepsilon} n)$, using for example the Schönhage–Strassen algorithm [SS71]. According to Fillebrown's analysis [Fil92, p. 441], the zeta-function algorithm for computing B_k requires $O(k)$ multiplications of integers with $O(k \log k)$ bits, so its running time is $O(k^2 \log^{2+\varepsilon} k)$. The new algorithm also runs in time $O(k^2 \log^{2+\varepsilon} k)$. However, we will see that over the feasible range of computation the running time behaves more like $O(k^2 \log k)$.

The new algorithm is easily parallelizable, as the computations for the various p are independent. Only the final stages of the modular reconstruction are more difficult to execute in parallel, but these account for only $O(k^{1+\varepsilon})$ of the running time. During the modular computations, each thread requires only $O(\log^{1+\varepsilon} k)$ space. The zeta-function algorithm may also be parallelized, for instance by mapping Euler factors to threads, but this requires $O(k \log k)$ space per thread.

2. COMPUTING B_k MODULO p

Throughout this section $p \geq 5$ denotes a prime, and k denotes an even integer satisfying $2 \leq k \leq p-3$. Our algorithm for computing B_k modulo p is based on the following well-known congruence. Let $0 < c < p$, and suppose that $c^k \not\equiv 1 \pmod{p}$. Then

$$(2) \quad B_k \equiv \frac{k}{1-c^k} \sum_{x=1}^{p-1} x^{k-1} h_c(x) \pmod{p},$$

where

$$h_c(x) := \frac{(x \bmod p) - c(x/c \bmod p)}{p} + \frac{c-1}{2}.$$

Equation (2) may be deduced by reading the statement of Theorem 2.3 of [Lan90, §2] modulo p (as is done in the proof of Corollary 2 to that theorem).

Equation (2) may be used to compute B_k modulo p directly, but computing x^{k-1} modulo p for each x is unnecessarily expensive. Instead, we select a generator g of the multiplicative group $(\mathbf{Z}/p\mathbf{Z})^\times$, put $c = g$, and rewrite the sum as

$$(3) \quad B_k \equiv \frac{k}{1-g^k} \sum_{i=1}^{p-1} g^{i(k-1)} h_g(g^i) \pmod{p}.$$

Note that $g^k \not\equiv 1 \pmod{p}$ since $p-1 \nmid k$. For $x \not\equiv 0 \pmod{p}$ we have $h_c(-x) = -h_c(x)$ and $g^{(p-1)/2} \equiv -1 \pmod{p}$, so we obtain the half-length sum

$$(4) \quad B_k \equiv \frac{2k}{1 - g^k} \sum_{i=1}^{(p-1)/2} g^{i(k-1)} h_g(g^i) \pmod{p},$$

leading to the following algorithm.

Algorithm 1: Compute B_k modulo p

```

 $g \leftarrow$  a generator of  $(\mathbf{Z}/p\mathbf{Z})^\times$ 
 $r \leftarrow g^{k-1} \pmod{p}$ 
 $u \leftarrow (g - 1)/2 \pmod{p}$ 
 $S \leftarrow 0, X \leftarrow 1, Y \leftarrow r$ 
for  $i \leftarrow 1$  to  $(p - 1)/2$  do
     $q \leftarrow \lfloor gX/p \rfloor$ 
     $S \leftarrow (S + (u - q)Y) \pmod{p}$ 
     $X \leftarrow gX \pmod{p}, Y \leftarrow rY \pmod{p}$ 
end
return  $2kS/(1 - g^k) \pmod{p}$ 

```

Proposition 1. *Let $p \geq 5$ be a prime, and let k be an even integer in the interval $2 \leq k \leq p - 3$. Algorithm 1 computes B_k modulo p in time $O(p \log^{1+\varepsilon} p)$.*

Proof. At the top of the loop we have $X = g^{i-1} \pmod{p}$ and $Y = g^{i(k-1)} \pmod{p}$. The value assigned to q is $(g(g^{i-1} \pmod{p}) - (g^i \pmod{p}))/p = u - h_g(g^i)$. Therefore the value added to S is $g^{i(k-1)} h_g(g^i)$. By (4) the return value is B_k modulo p .

We now analyze the complexity. A generator of $(\mathbf{Z}/p\mathbf{Z})^\times$ may be found deterministically in time $O(p^{1/4+\varepsilon})$ [Shp96]. In the main loop, the algorithm performs $O(p)$ additions, multiplications, and divisions (with remainder) of integers with $O(\log p)$ bits. The computation of r requires another $O(\log k) = O(\log p)$ such operations. The division by $1 - g^k$ requires an extended GCD of integers with $O(\log p)$ bits. The total cost is $O(p \log^{1+\varepsilon} p)$. \square

Remark. To compute B_{10^8} (see §5), the largest prime needed is 1,558,322,063. Even on a 32-bit machine, this fits into a single machine word, and the cost of arithmetic modulo p does not depend on p . Therefore, over the feasible range of computation, the complexity of Algorithm 1 may be regarded as only $O(p)$.

3. COMPUTING B_k AS A RATIONAL NUMBER

In this section we present a multimodular algorithm for computing $B_k = N_k/D_k$ as an exact rational number. Before getting to the algorithm proper, we make some preliminary calculations. In what follows, we assume that $k \geq 4$ and that k is even.

The denominator D_k is given by the von Staudt–Clausen theorem [IR90, p. 233]:

$$D_k = \prod_{\substack{p \text{ prime} \\ p-1|k}} p.$$

Note that $D_k \leq 2^{k+1}$, since D_k divides $2(2^k - 1)$ [CH72, p. 114].

The size of the numerator N_k may be bounded quite tightly as follows. From (1) and Stirling's approximation [Lan97, p. 120] we have

$$\begin{aligned} \log |B_k| &= \log 2 + \log \zeta(k) + \log k! - k \log 2\pi \\ &\leq \left(k + \frac{1}{2}\right) \log k - (1 + \log 2\pi)k + (\log 2 + \log \zeta(k) + \log \sqrt{2\pi}) + \frac{1}{12k}. \end{aligned}$$

Since $\zeta(k) \leq \zeta(4)$ and $12k \geq 48$, a short calculation shows that $|N_k| < 2^\beta$ where $\beta := \lceil (k + 0.5) \log_2 k - 4.094k + 2.470 + \log_2 D_k \rceil$.

For any $X > 0$, let $M_X := \prod_{p \leq X, p-1 \nmid k} p$. Our strategy will be to select X so that $M_X \geq 2^\beta$ and then compute N_k modulo M_X . This ensures that we have sufficient information to reconstruct N_k . In the algorithm itself our choice of X will be quite sharp, but we also need a rough *a priori* estimate. We will take $Y := \max(37, \lceil (k + 0.5) \log_2 k \rceil)$. To check that this works, we will use the estimate $\sum_{p \leq x} \log p \geq 0.73x$, a weak form of the prime number theorem, valid for $x \geq 37$ [Nar00, p. 111]. Then we have

$$\begin{aligned} \log_2 M_Y &= -\log_2 D_k + \sum_{p \leq Y} \log_2 p \\ &\geq -(k+1) + \frac{0.73}{\log 2} Y \\ &\geq Y - k - 1 \\ &\geq (k + 0.5) \log_2 k - k - 1 \\ &\geq (k + 0.5) \log_2 k - 3.094k + 6.470 \quad (\text{since } k \geq 4) \\ &\geq (k + 0.5) \log_2 k - 4.094k + 3.470 + \log_2 D_k + 2 \\ &\geq \lceil (k + 0.5) \log_2 k - 4.094k + 2.470 + \log_2 D_k \rceil + 2 \\ &= \beta + 2, \end{aligned}$$

so that $M_Y \geq 2^{\beta+2}$. Algorithm 2 presents the resulting algorithm; several important details are given in the proof of Theorem 2.

Theorem 2. *Algorithm 2 computes B_k in time $O(k^2 \log^{2+\varepsilon} k)$.*

Proof. Computing all primes less than Y requires time $O(k^{1+\varepsilon})$, via an elementary sieving algorithm, since $Y = O(k \log k)$. In what follows, we assume that all primality tests and enumerations of primes are performed with the aid of this list.

To compute D_k , make a list of the factors of k (for example, by testing divisibility of k by each integer in $[2, \sqrt{k}]$), and for each factor $d|k$ check whether $d+1$ is prime. Multiply together those that are prime using a product tree [vzGG03, Algorithm 10.3, p. 293]. Since $\log D_k = O(k)$, this takes time $O(k^{1+\varepsilon})$.

The goal of the while-loop is to quickly find a bound X , much tighter than Y , such that $M_X \geq 2^\beta$. The variable M' holds an approximation to the product of the primes encountered so far. It should be represented in floating-point, with at least $n = \lceil \log_2 Y \rceil + 1$ bits in the mantissa, so that each multiplication by p magnifies the relative error by at most $1 + 2^{-n}$. We claim that the loop terminates before exhausting the precomputed list of primes. Indeed, let s be the number of primes p such that $p \leq Y$ and $p-1 \nmid k$. After s iterations, M' approximates M_Y with relative error at most $(1 + 2^{-n})^s \leq (1 + (2Y)^{-1})^Y \leq e^{1/2}$; that is, $M' \geq e^{-1/2} M_Y \geq e^{-1/2} 2^{\beta+2} \geq 2^{\beta+1}$. This is impossible since the loop termination condition would

Algorithm 2: Compute B_k as a rational number ($k \geq 4$, k even)

```

// compute  $D_k$  and preliminary bounds
 $Y \leftarrow \max(37, \lceil (k + 0.5) \log_2 k \rceil)$ 
Compute list of all primes  $p \leq Y$ 
 $D_k \leftarrow \prod_{p-1|k} p$ 
 $\beta \leftarrow \lceil (k + 0.5) \log_2 k - 4.094k + 2.470 + \log_2 D_k \rceil$ 
// compute tight bound  $X$ 
 $p \leftarrow 3, M' \leftarrow 1.0$ 
while  $M' < 2^{\beta+1}$  do
     $p \leftarrow \text{NextPrime}(p)$ 
    if  $p - 1 \nmid k$  then  $M' \leftarrow pM'$ 
end
 $X \leftarrow p$ 
// collect modular information
for  $p \in \{2, 3, 5, \dots, X\}, p - 1 \nmid k$  do  $r_p \leftarrow B_k \pmod{p}$ 
// reconstruction
 $M \leftarrow M_X = \prod_{p \leq X, p-1 \nmid k} p$ 
 $R \leftarrow$  unique integer in  $[0, M)$  congruent to  $r_p$  modulo  $p$  for all primes  $p \mid M$ 
 $N' \leftarrow D_k R \pmod{M}$ 
if  $k \equiv 2 \pmod{4}$  then  $N_k \leftarrow N'$  else  $N_k \leftarrow N' - M$ 
return  $N_k/D_k$ 

```

already have been met earlier. This argument also shows that the selected bound X satisfies $M_X \geq e^{-1/2} 2^{\beta+1} \geq 2^\beta$. Since $Y = O(k \log k)$, computing X takes time $O(k^{1+\epsilon})$.

In the for-loop, for each p we use Algorithm 1 to compute $r_p := B_k \pmod{p}$. Note that Algorithm 1 requires that $2 \leq k \leq p - 3$. To satisfy this requirement, we put $m = k \pmod{p - 1}$, compute $B_m \pmod{p}$ using Algorithm 1, and then recover $B_k \pmod{p}$ via Kummer's congruence $B_k/k \equiv B_m/m \pmod{p}$ [Lan90, p. 41]. The total number of primes processed is no more than $\pi(Y) = O(Y/\log Y) = O(k \log k / \log(k \log k)) = O(k)$. According to Proposition 1, the cost for each prime is $O(p \log^{1+\epsilon} p) = O(k \log^{2+\epsilon} k)$, since $p \leq Y = O(k \log k)$. Therefore the total cost of the while-loop is $O(k^2 \log^{2+\epsilon} k)$.

We compute M and R simultaneously using fast Chinese Remaindering [vzGG03, Theorem 10.25, p. 302]; the cost is $O((\log M_X)^{1+\epsilon}) = O((k \log k)^{1+\epsilon}) = O(k^{1+\epsilon})$.

In the last few lines we recover N_k . By construction we have $R \equiv B_k \pmod{M}$, so $N' \equiv N_k \pmod{M}$. Note that $|N_k| < 2^\beta \leq M$. If $k \equiv 2 \pmod{4}$, then N_k is positive, so we simply have $N_k = N'$; otherwise N_k is negative, and we must correct N' by subtracting M . □

Remark. During the modular collection phase, we skipped those p for which $p - 1 \mid k$. An alternative would be to use the fact that $pB_k \equiv -1 \pmod{p}$ for these p [IR90, p. 233] and to apply the multimodular algorithm directly to N_k rather than to B_k . This would require the additional minor overhead of computing $D_k \pmod{p}$ for all of the other primes.

Remark. As mentioned below the proof of Proposition 1, the running time of Algorithm 1 behaves like $O(p)$ in practice. The corresponding bound for the running time of Algorithm 2 is $O(k^2 \log k)$.

4. IMPLEMENTATION TECHNIQUES FOR COMPUTING B_k MODULO p

In this section we sketch several techniques that yield a large constant factor improvement in the speed of computing B_k modulo p , for almost all k and p . In experiments we have found that these techniques yield an increase in speed by a factor in excess of 150 compared to an efficient implementation of Algorithm 1.

We first perform some preparatory algebra. Consider again the congruence (2). Let n be the multiplicative order of c in $(\mathbf{Z}/p\mathbf{Z})^\times$, and put $m = (p - 1)/n$. Let g be a generator of $(\mathbf{Z}/p\mathbf{Z})^\times$. Then $\{g^i c^{-j}\}_{0 \leq i < m, 0 \leq j < n}$ forms a complete set of representatives for $(\mathbf{Z}/p\mathbf{Z})^\times$. Putting $x = g^i c^{-j}$, we obtain

$$(5) \quad B_k \equiv \frac{k}{1 - c^k} \sum_{\substack{0 \leq i < m \\ 0 \leq j < n}} (g^{k-1})^i (c^{k-1})^{-j} h_c(g^i c^{-j}) \pmod{p}.$$

Considerations similar to those of §2 allow us to halve the number of terms, as follows. First suppose that n is even. Since $h_c(-x) = -h_c(x)$ and $c^{n/2} \equiv -1 \pmod{p}$, the double sum in (5) becomes

$$\begin{aligned} & \sum_{\substack{0 \leq i < m \\ 0 \leq j < n/2}} (g^{k-1})^i (c^{k-1})^{-j} h_c(g^i c^{-j}) + \sum_{\substack{0 \leq i < m \\ 0 \leq j < n/2}} (g^{k-1})^i (c^{k-1})^{-j-n/2} h_c(g^i c^{-j-n/2}) \\ \equiv & 2 \sum_{\substack{0 \leq i < m \\ 0 \leq j < n/2}} (g^{k-1})^i (c^{k-1})^{-j} h_c(g^i c^{-j}) \pmod{p}. \end{aligned}$$

Now suppose that n is odd. Then m is even, and we have $(-g^{m/2})^n \equiv (-1)^n g^{(p-1)/2} \equiv (-1)(-1) \equiv 1 \pmod{p}$, so that $-g^{m/2} \equiv c^\lambda \pmod{p}$ for some $0 \leq \lambda < n$. The double sum in (5) becomes

$$\begin{aligned} & \sum_{\substack{0 \leq i < m/2 \\ 0 \leq j < n}} (g^{k-1})^i (c^{k-1})^{-j} h_c(g^i c^{-j}) + \sum_{\substack{0 \leq i < m/2 \\ 0 \leq j < n}} (g^{k-1})^{i+m/2} (c^{k-1})^{-j} h_c(g^{i+m/2} c^{-j}) \\ \equiv & \sum_{\substack{0 \leq i < m/2 \\ 0 \leq j < n}} (g^{k-1})^i (c^{k-1})^{-j} h_c(g^i c^{-j}) - \sum_{\substack{0 \leq i < m/2 \\ 0 \leq j < n}} (g^{k-1})^i (c^{k-1})^{-j+\lambda} h_c(-g^i c^{-j+\lambda}) \\ \equiv & 2 \sum_{\substack{0 \leq i < m/2 \\ 0 \leq j < n}} (g^{k-1})^i (c^{k-1})^{-j} h_c(g^i c^{-j}) \pmod{p}. \end{aligned}$$

We combine both cases into a single statement by writing

$$(6) \quad B_k \equiv \frac{2k}{1 - c^k} \sum_{0 \leq i < m'} (g^{k-1})^i \sum_{0 \leq j < n'} (c^{k-1})^{-j} h_c(g^i c^{-j}) \pmod{p},$$

where

$$n' = \begin{cases} n/2, & n \text{ even,} \\ n, & n \text{ odd,} \end{cases} \quad m' = \frac{(p-1)/2}{n'}.$$

The sum in (6) may be evaluated by an algorithm similar to Algorithm 1, with an outer loop over i and an inner loop over j ; we omit the details. The inner loop is executed $n'm' = (p - 1)/2$ times, the same number of iterations as in the main loop of Algorithm 1. Note that if c is chosen ‘randomly’, then we expect n' to be quite large, so it is imperative to make the inner loop as efficient as possible.

To this end, we specialize to the case $c = 1/2 \pmod p$; this will allow us to exploit the fact that multiplication by 2 modulo p is very cheap. Of course, this specialization does not work if $2^k \equiv 1 \pmod p$. In practice, such primes tend to be rare for any given k , and we may fall back on Algorithm 1 to handle them.

Directly from the definition of $h_c(x)$, we have $h_{1/2}(x) = -f(x)/4$ where

$$f(x) := \begin{cases} 1, & 0 < (x \pmod p) < p/2, \\ -1, & p/2 < (x \pmod p) < p. \end{cases}$$

From (6) we obtain

$$(7) \quad B_k \equiv \frac{k}{2(2^{-k} - 1)} \sum_{0 \leq i < m'} (g^{k-1})^i \sum_{0 \leq j < n'} (2^{k-1})^j f(g^i 2^j) \pmod p,$$

where n is the order of 2 in $(\mathbf{Z}/p\mathbf{Z})^\times$ and n' and m' are defined as before.

In the algorithm that evaluates (7), the inner loop is considerably simpler than the main loop of Algorithm 1. Given $g^i 2^j$ at the beginning of the loop, we compute the next term $g^i 2^{j+1}$ by doubling it and subtracting p if necessary, and $f(g^i 2^j)$ is -1 or $+1$ according to whether the subtraction occurred. Therefore only a single modular multiplication is required on each iteration, to keep track of $(2^{k-1})^j$.

Next we describe further optimizations for computing a sum of the form

$$(8) \quad \sum_{0 \leq j < N} r^j f(2^j s) \pmod p.$$

(For evaluating (7), we take $N = n'$, $r = 2^{k-1}$, $s = g^i$.) The following observation is crucial: if $0 < s < p$ and if the binary expansion of s/p is $0.b_0 b_1 b_2 \dots$ (where each b_j is 0 or 1), then $f(2^j s) = (-1)^{b_j}$. Indeed, the binary expansion of $2^j s/p$ is obtained by shifting that of s/p to the left by j digits, so

$$(2^j s \pmod p)/p = 0.b_j b_{j+1} \dots,$$

and then

$$f(2^j s) = +1 \iff 0 < (2^j s \pmod p) < p/2 \iff b_j = 0.$$

Assume that we have available a routine for computing the binary expansion of s/p , whose output is a sequence of base-2^m digits. Strip off the last $N \pmod m$ terms of (8) and compute them separately; we may then assume that N is divisible by m . Put $j = mt + u$, so that (8) becomes

$$(9) \quad \sum_{0 \leq t < N/m} (r^m)^t \sum_{0 \leq u < m} r^u f(2^{mt+u} s).$$

We may now use a caching strategy as follows. Maintain a table $\{T_\sigma\}$ of length 2^m , where $\sigma \in \{+1, -1\}^m$. Each T_σ is a residue modulo p and is initialized to zero. For each $0 \leq t < N/m$, add $(r^m)^t$ to the table element T_σ , where σ is determined from the t -th base-2^m digit of s/p by

$$\sigma = (f(2^{mt} s), f(2^{mt+1} s), \dots, f(2^{mt+m-1} s)).$$

(In practice, σ is represented by the sequence of bits themselves and is obtained by a simple bit-mask.) After finishing the loop over t , compute the 2^m values $V_\sigma = \sum_{0 \leq u < m} r^u \sigma_u$, where $\sigma = (\sigma_0, \dots, \sigma_{m-1})$. Then the desired sum is given by $\sum_\sigma V_\sigma T_\sigma$. The main benefit of this approach is that in the inner loop we only perform N/m modular multiplications, instead of N . The tradeoff is that we must maintain the table $\{T_\sigma\}$, and some extra work is required at the end to assemble the

information from the table. For this to be worthwhile, we should have $N \gg 2^m$. We should also choose m so that the base- 2^m digits of s/p can be extracted efficiently. Moreover, memory locality problems can offset the savings in arithmetic if m is too large. We have found in practice that $m = 8$ works quite well.

In the implementation discussed in §5 below, we made several further optimizations, which we describe only briefly.

For better memory locality — and indeed less total memory consumption — we split (9) into blocks and process each block separately. This avoids needing to store and then later retrieve too many bits of the expansion of s/p . Instead of computing s/p for each s separately, we precompute an approximation to $1/p$ and then multiply it by each s that occurs, taking advantage of the speed of multiplication by a single-word integer compared to the corresponding division.

We use several tables instead of just one. For example, on a 32-bit machine, we use four tables: the highest 8 bits of each word of s/p contribute to the first table, the next 8 bits to the second table, and so on. This reduces further the number of modular multiplications. When adding values into a table, we skip the reduction modulo p , delaying this until the end of the loop. Of course, this is only possible when the word size is large enough compared to p ; otherwise overflow may occur. Finally, instead of reinitializing the table on each iteration of the outer loop, we simply continue to accumulate data into the same table. This partly mitigates against the overhead incurred when m' is unusually large compared to n' .

Most of the modular arithmetic uses Montgomery's reduction algorithm [Mon85].

5. PERFORMANCE DATA

The author implemented the above algorithms in a C++ package named `bernm` (*Bernoulli multi-modular*). The source code is available from the author's web page under a free software license. It depends on the GMP library [Gra08] for arbitrary precision arithmetic and on NTL [Sho07] for some of the single-precision modular arithmetic. It is included as a standard component of the Sage computer algebra system (version 3.1.2) [SJ05], accessible via the `bernoulli` function.

The parallel version distributes the modular computations among the requested number of threads and depends on the standard pthreads API. For the reconstruction phase, the lowest layers of the subproduct tree are performed in parallel, but the number of threads decreases towards the top of the tree, since the extended GCD routine is not threaded.

Table 1 shows timing data for $10^4 \leq k \leq 10^8$, spaced by intervals of $\sqrt{10}$, for `bernm` and several existing implementations of the zeta-function algorithm. The timings were obtained on a 16-core 2.6GHz AMD Opteron (64-bit) machine with 96 GB RAM, running Ubuntu Linux. The last two rows of the table ($k = \lfloor 10^{7.5} \rfloor$ and $k = 10^8$) improve on the prior record $k = 10^7$ set by Pavlyk [Pav08]. The values $B_{31622776}$ and B_{10^8} may be downloaded from the author's web page.

The compiler used was gcc 4.1.3. We used the version of GMP shipped with Sage 3.1.2, which is the same as GMP 4.2.2 but also includes assembly code written by Pierrick Gaudry that improves performance on the Opteron and code by Niels Möller that implements a quasi-linear time extended GCD algorithm [Möl08]. We used NTL version 5.4.1, also included in Sage.

The timings for PARI/GP were obtained for version 2.3.3, using the same patched version of GMP. We used the `bernfrac` function from the GP interpreter. The

TABLE 1. Time (in seconds) for computing B_k for PARI/GP, Mathematica, calcbn, and bernmm.

k	threads	PARI	MMA	calcbn		bernmm	
		CPU time	CPU time	CPU time	wall time	CPU time	wall time
10^4	1	0.07	0.18	0.04		0.25	
$[10^{4.5}]$	1	0.68	1.80	0.49		1.02	
10^5	1	7.45	21.7	5.72		5.46	
$[10^{5.5}]$	1	88.3	285	82.2		38.6	
	2			82.3	41.6	38.6	20.4
10^6	1	1058	3434	935		340	
	4			967	246	341	96.1
$[10^{6.5}]$	1	$1.24 \cdot 10^4$	$4.25 \cdot 10^4$			3397	
	4			$1.33 \cdot 10^4$	3385	3406	909
10^7	1	$1.42 \cdot 10^5$	$5.10 \cdot 10^5$				
	10			$1.60 \cdot 10^5$	$1.61 \cdot 10^4$	$3.59 \cdot 10^4$	3938
$[10^{7.5}]$	10					$3.86 \cdot 10^5$	$4.03 \cdot 10^4$
10^8	10					$4.19 \cdot 10^6$	$4.27 \cdot 10^5$

timings for calcbn, a specialized program for computing Bernoulli numbers, were obtained from calcbn 2.0 [Kel08], again using the same patched GMP. The timings for Mathematica were obtained from the Linux x86 (64-bit) version of Mathematica 6.0, using the `BernoulliB` function. Neither PARI/GP nor Mathematica supply a parallel implementation as far as we are aware.

The peak memory usage for computing B_{10^8} was approximately 5 GB; this occurred during the final extended GCD operation. By comparison, PARI/GP used 4.3 GB for $k = 10^7$, and calcbn used 1.7 GB for $k = 10^7$ (with ten threads).

For each implementation, we observe that the ratio of the times in adjacent rows of the table is approximately 10, reflecting the predicted quasi-quadratic asymptotic running time of both the zeta-function algorithm and the multimodular algorithm.

We checked our results by two methods. First, we checked that as a real number, the proposed value for B_k agrees with the estimate $|B_k| \approx 2(2\pi)^{-k}k!$. Since we computed B_k by modular information alone, this is a very strong consistency check. Second, we reduced modulo p the proposed value for B_k and compared this against the output of Algorithm 1, for a few primes *distinct* from those primes that we used to compute B_k . Again, this is a very strong consistency check — stronger in fact, since it involves all of the bits of the proposed B_k .

Mathematica, PARI/GP, and calcbn use GMP internally, so naturally any improvement in GMP’s large integer arithmetic will improve the timings reported in the table. In this connection we mention the improvements in multiplication speed reported by [GKZ07]; their paper also gives comparisons with other implementations of large-integer arithmetic.

ACKNOWLEDGMENTS

Many thanks to Joe Buhler, Bernd Kellner, and Andrew Sutherland for their comments on a draft of this paper and to the Department of Mathematics at Harvard University for providing the hardware on which the computations were performed.

REFERENCES

- [CH72] S. Chowla and P. Hartung, *An “exact” formula for the m -th Bernoulli number*, Acta Arith. **22** (1972), 113–115. MR0308036 (46:7151)
- [Fil92] Sandra Fillebrown, *Faster computation of Bernoulli numbers*, J. Algorithms **13** (1992), no. 3, 431–445. MR1176671 (94d:68044)
- [FP07] Greg Fee and Simon Plouffe, *An efficient algorithm for the computation of Bernoulli numbers*, 2007, preprint <http://arxiv.org/abs/math/0702300>.
- [GKZ07] Pierrick Gaudry, Alexander Kruppa, and Paul Zimmermann, *A GMP-based Implementation of Schönhage–Strassen’s Large Integer Multiplication Algorithm*, ISSAC (Dongming Wang, ed.), ACM, 2007, pp. 167–174. MR2396199 (2009e:11236)
- [Gra08] Torbjörn Granlund, *The GNU Multiple Precision Arithmetic library*, 2008, <http://gmp.lib.org/>.
- [IR90] Kenneth Ireland and Michael Rosen, *A classical introduction to modern number theory*, second ed., Graduate Texts in Mathematics, vol. 84, Springer-Verlag, 1990. MR1070716 (92e:11001)
- [Kel02] Bernd C. Kellner, *Über irreguläre Paare höherer Ordnungen*, Diplomarbeit, 2002.
- [Kel08] ———, *calcbn*, 2008, <http://www.bernoulli.org/>.
- [Lan90] Serge Lang, *Cyclotomic fields I and II*, second ed., Graduate Texts in Mathematics, vol. 121, Springer-Verlag, New York, 1990. MR1029028 (91c:11001)
- [Lan97] ———, *Undergraduate analysis*, second ed., Undergraduate Texts in Mathematics, Springer-Verlag, New York, 1997. MR1476913 (98h:00002)
- [Möl08] Niels Möller, *On Schönhage’s algorithm and subquadratic integer GCD computation*, Math. Comp. **77** (2008), no. 261, 589–607 (electronic). MR2353968 (2008h:11004)
- [Mon85] Peter L. Montgomery, *Modular multiplication without trial division*, Math. Comp. **44** (1985), no. 170, 519–521. MR777282 (86e:11121)
- [Nar00] Władysław Narkiewicz, *The development of prime number theory*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2000, From Euclid to Hardy and Littlewood. MR1756780 (2001c:11098)
- [Par07] The PARI Group, Bordeaux, *PARI/GP, version 2.3.3*, 2007, available from <http://pari.math.u-bordeaux.fr/>.
- [Pav08] Oleksandr Pavlyk, *Today We Broke the Bernoulli Record: From the Analytical Engine to Mathematica*, 2008, posted on Wolfram Blog (<http://blog.wolfram.com/>) on 29th April 2008, retrieved 15th June 2008.
- [Sho07] Victor Shoup, *NTL: A library for doing number theory*, <http://www.shoup.net/ntl/>, 2007.
- [Shp96] Igor Shparlinski, *On finding primitive roots in finite fields*, Theoret. Comput. Sci. **157** (1996), no. 2, 273–275. MR1389773 (97a:11203)
- [SJ05] William Stein and David Joyner, *Sage: System for algebra and geometry experimentation*, Communications in Computer Algebra (ACM SIGSAM Bulletin) **39** (2005), no. 2, 61–64.
- [SS71] A. Schönhage and V. Strassen, *Schnelle Multiplikation grosser Zahlen*, Computing (Arch. Elektron. Rechnen) **7** (1971), 281–292. MR0292344 (45:1431)
- [Ste07] William Stein, *Modular forms, a computational approach*, Graduate Studies in Mathematics, vol. 79, American Mathematical Society, Providence, RI, 2007, with an appendix by Paul E. Gunnells. MR2289048 (2008d:11037)
- [vzGG03] Joachim von zur Gathen and Jürgen Gerhard, *Modern computer algebra*, second ed., Cambridge University Press, Cambridge, 2003. MR2001757 (2004g:68202)
- [Wol07] Wolfram Research, Inc., *Mathematica 6.0*, 2007, <http://www.wolfram.com/>.

COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY, 251 MERCER STREET, NEW YORK, NEW YORK 10012

E-mail address: dmharvey@cims.nyu.edu