Schönhage-Strassen Algorithm with MapReduce for Multiplying Terabit Integers (April 29, 2011)

Tsz-Wo Sze
Yahoo! Cloud Platform
701 First Avenue
Sunnyvale, CA 94089, USA
tsz@yahoo-inc.com

ABSTRACT

We present *MapReduce-SSA*, an integer multiplication algorithm using the ideas from Schönhage-Strassen algorithm (SSA) on MapReduce. SSA is one of the most commonly used large integer multiplication algorithms. MapReduce is a programming model invented for distributed data processing on large clusters. MapReduce-SSA is designed for multiplying integers in terabit scale on clusters of commodity machines. As parts of MapReduce-SSA, two algorithms, *MapReduce-FFT* and *MapReduce-Sum*, are created for computing discrete Fourier transforms and summations. These mathematical algorithms match the model of MapReduce seamlessly.

Categories and Subject Descriptors

I.1.2 [**Computing Methodologies**]: Algorithms—*symbolic and algebraic manipulation*; F.2.1 [**Theory of Computation**]: Numerical Algorithms—*computation of fast Fourier transforms*; G.4 [**Mathematics of Computing**]: Mathematical Software

General Terms

Algorithms, distributed computing

Keywords

Integer multiplication, multiprecision arithmetic, fast Fourier transform, summation, MapReduce

1. INTRODUCTION

Integer multiplication is one of the most critical operations in arbitrary precision arithmetic. Beyond its direct applications, the existence of fast multiplication algorithms running in essentially linear time and the reducibility of other common operations such as division, square root, logarithm, etc. to integer multiplication motivate improvements in techniques for its efficient computation at scale [4, 8]. For ter-

abit scale integer multiplication, the major application of it is the computation of classical constants, in particular, the computation of π [21, 1, 18, 7, 5].

The known, asymptotically fastest multiplication algorithms are based on fast Fourier transform (FFT). In practice, multiplication libraries often employ a mix of algorithms, as the performance and suitability of an algorithm vary depending on the bit range computed and on the particular machine where the multiplication is performed. While benchmarks inform the tuning of an implementation on a given computation platform, algorithms best suited to particular bit ranges are roughly ranked both by empirical and theoretical results. Generally, the naïve algorithm performs best in the lowest bit range, then the Karatsuba algorithm, then the Toom-Cook algorithm, and multiplication in the largest bit ranges typically employs FFT-based algorithms. For details, see [3, 13].

The Schönhage-Strassen algorithm (SSA) is one such FFT-based integer multiplication algorithm, requiring

$$O(N \log N \log \log N)$$

bit operations to multiply two N-bit integers [15]. There are other FFT-based algorithms asymptotically faster than SSA. In 2007, Fürer published an algorithm with running time

$$O(N(\log N)2^{\log^* N})$$

bit operations [11]. Using similar ideas from Fürer's algorithm, De et al. discovered another $O(N(\log N)2^{\log^* N})$ algorithm using modular arithmetic while the arithmetic of Fürer's algorithm is carried out over complex numbers [9]. As these two algorithms are relatively new, they are not found in common software packages and the bit ranges where they begin to outperform SSA are, as yet, unclear. SSA remains a dominant implementation in practice. See [12, 6] for details on SSA implementations.

Algorithms for multiplying large integers, including SSA, require a large amount of memory to store the input operands, intermediate results, and output product. For in-place SSA, the total space required to multiply to N-bit integers ranges from $8N$ to $10N$ bits; see §3.1. Traditionally, multiplication of terabit-scale integers is performed on supercomputers with software customized to exploit its particular memory and interconnect architecture [5, 7, 18]. Takahashi recently computed the square of a 2 TB integer with 5.15 trillion decimal digits in 31 minutes and 41 seconds using the 17.5 TB main memory on the T2K Open Supercomputer [17].

Most arbitrary precision arithmetic software packages assume a uniform address space on a single machine. The

precision supported by libraries such as GMP and Magma is limited by the available physical memory. When memory is exhausted during a computation, these packages may have very poor performance or may not work at all. Such packages are not suited to terabit scale multiplication, as the memory available in most machines is insufficient by orders of magnitude. Other implementations such as *y-cruncher* [20] and *TachusPI* [2] use clever paging techniques to scale a single machine beyond these limits. By using the hard disk sparingly, its significantly higher capacity enables efficient computation despite its high latencies. According to Yee, y-cruncher is able to square an integer with 5.2 trillion decimal digits using a desktop computer with only 96 GB memory in 41.6 hours [19].

In this paper, we present a distributed variety of SSA for multiplying terabit integers on commodity hardware using a shared compute layer. Our prototype implementation runs on *Hadoop* clusters. Hadoop is an open source, distributed computing framework developed by the Apache Software Foundation. It includes a fault-tolerant, distributed file system, namely *HDFS*, designed for high-throughput access to very large data sets in petabyte scale [16]. It also includes an implementation of *MapReduce*, a programming model designed for processing big data sets on large clusters [10].

Our environment is distinguished both from the supercomputer and single machine variants described in the preceding paragraphs. Unlike these settings, our execution environment assumes that hardware failures are common. The fixed computation plan offered by our host environment is designed for multi-tenant, concurrent batch processing of large data volumes on a shared storage and compute grid. As a sample constraint, communication patterns between processes are fixed and communicated only at their conclusion, to ensure execution plans remain feasible should a machine fail. While the problem of memory remains a central and familiar consideration of the implementation, some of the challenges posed and the advantages revealed by our programming paradigm and execution environment are unique. We learned that not only SSA, but FFT and other arbitrary precision routines like summation match the MapReduce programming model seamlessly. More specifically, the FFT matrix transposition, which is traditionally difficult in preserving locality, becomes trivial in MapReduce.

The rest of the sections are organized as follows. More details on computation environment are given in §2. SSA is briefly described in §3. A new algorithm, *MapReduce-SSA*, is presented in §4. The conclusion is in §5.

2. COMPUTATION ENVIRONMENT

In this section, we first describe our cluster configurations, and then give a short introduction of MapReduce.

2.1 Cluster Configurations

At Yahoo!, Hadoop clusters are composed of hundreds or thousands of commodity machines. All the machines in the same cluster are collocated in the same data center. Each machine typically has

- 2 quad-core CPUs,

- 6 GB or more memory, and

- 4 hard drives.

The machines are connected with a two-level network topology shown in Figure 1. There are 40 machines connected to a rack switch by 1-gigabit links. Each rack switch is connected to a set of core switches with 8-gigabit or higher aggregate bandwidth.
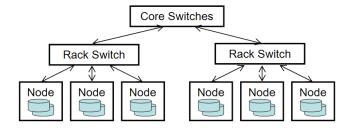

Figure 1: Network topology

For terabit scale integer multiplication, the data involved in the matrix transposition step in parallel FFT is terabyte scale. Small clusters with tens of machines are inadequate to perform such operation, even if individual machines are equipped with high performance CPUs and storage, since the network bandwidth is a bottleneck of the system. The benefits of adding machines to a cluster are twofold: it increases the number of CPUs and the size of storage, and also increases the aggregate network bandwidth. From our experience, clusters with five hundred machines are sufficient to perform terabit scale multiplication.

2.2 MapReduce

In the MapReduce model, the input and the output of a computation are lists of key-value pairs. Let k_t be key types and v_t be value types for $t = 1, 2, 3$. The user specifies two functions,

$$\text{map}: \quad (k_1, v_1) \longrightarrow \text{list}\langle k_2, v_2\rangle, \quad \text{and}$$
$$\text{reduce}: \quad (k_2, \text{list}\langle v_2\rangle) \longrightarrow \text{list}\langle k_3, v_3\rangle.$$

When a *job* starts, the MapReduce framework launches *map tasks*. A map task transforms one or more inputs in (k_1, v_1) to intermediate outputs in (k_2, v_2) according to map($\cdot$). Once all map tasks have completed, the framework begins a process called *shuffle*. It puts all intermediate values with the same key into a list, launches *reduce tasks* and sends the keys with the corresponding list to the reduce tasks. A reduce task uses reduce($\cdot$) to process one or more of the key-list pairs in $(k_2, \text{list}\langle v_2\rangle)$ and outputs key-value pairs in (k_3, v_3). The job is finished when all reduce tasks are completed. Figure 2 shows a diagram of the MapReduce model.
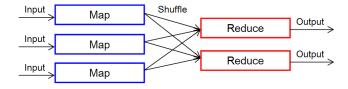

Figure 2: MapReduce model

The number of map tasks is at least one but the number of reduce tasks could possibly be zero. For a map-only job,

shuffle is unnecessary and is not performed. The output from the map tasks becomes the output of the job.

The MapReduce framework included in Hadoop is highly parallel, locality-aware, elastic, flexible and fault tolerant. Operators routinely add new machines and decommission failing ones without affecting running jobs. Hadoop provides a flexible API and permits users to define various components, including input/output formats that define how data is read and written from HDFS. Hardware failures are common, so the framework must also detect and gracefully tolerate such events. As a MapReduce application, recovery is transparent as long as it conforms to the assumptions and conventions of the framework.

3. SCHÖNHAGE-STRASSEN

We describe SSA briefly in this section. For more detailed discussions, see [15, 8, 3].

The goal is to compute

$$p = ab, \tag{1}$$

where $a, b \in \mathbb{Z}$ are the input integers and p is the product of a and b. Without loss of generality, we may assume $a > 0$ and $b > 0$. Suppose a and b are N-bit integers for N a power of two. Write

$$N = KM = 2^k M,$$

where K and M are also powers of two. Let

$$A(x) \stackrel{\text{def}}{=} \sum_{i=0}^{K-1} a_i x^i \qquad \text{for } 0 \le a_i < 2^M \qquad \text{and}$$

$$B(x) \stackrel{\text{def}}{=} \sum_{i=0}^{K-1} b_i x^i \qquad \text{for } 0 \le b_i < 2^M$$

be polynomials such that $A(x), B(x) \in \mathbb{Z}[x]$, $a = A(2^M)$ and $b = B(2^M)$. There are M bits in each coefficient a_i and b_i. The number of coefficients in each polynomial $A(x)$ and $B(x)$ is K. Let

$$D \stackrel{\text{def}}{=} 2K = 2^{k+1}.$$

The strategy is to compute the product polynomial

$$P(x) \stackrel{\text{def}}{=} A(x)B(x) \stackrel{\text{def}}{=} \sum_{i=0}^{D-1} p_i x^i \qquad \text{for } 0 \le p_i < 2^{2M+k}.$$

Then, we have $p = P(2^M)$. Consider polynomials A, B and P as D-dimensional tuples,

$$\begin{aligned} \mathbf{a} &\stackrel{\text{def}}{=} (0, \ldots, 0, a_{K-1}, \ldots, a_0), \\ \mathbf{b} &\stackrel{\text{def}}{=} (0, \ldots, 0, b_{K-1}, \ldots, b_0), \qquad \text{and} \\ \mathbf{p} &\stackrel{\text{def}}{=} (p_{D-1}, \ldots, p_0). \end{aligned}$$

Then $\mathbf{p}$ is the cyclic convolution of $\mathbf{a}$ and $\mathbf{b}$,

$$\mathbf{p} = \mathbf{a} \times \mathbf{b}.$$

By the convolution theorem,

$$\mathbf{a} \times \mathbf{b} = \text{dft}^{-1}(\text{dft}(\mathbf{a}) * \text{dft}(\mathbf{b})), \tag{2}$$

where $*$ denotes componentwise multiplication, $\text{dft}(\cdot)$ and $\text{dft}^{-1}(\cdot)$ respectively denote the discrete Fourier transform and its inverse. All the operations on the right hand side of equation (2) are performed over the ring $\mathbb{Z}/(2^n + 1)\mathbb{Z}$. The

integer $2^n + 1$ is called *the Schönhage-Strassen modulus*. The exponent n must satisfy the following conditions,

(C_1) $D \mid 2n$, and

(C_2) $n \ge 2M + k$.

Thus, n is chosen to be the smallest integer satisfying both conditions (C_1) and (C_2). At last, the componentwise multiplications are computed by recursive calls to the integer multiplication routine, and dft (or dft^{-1}) is calculated by forward (or backward) FFT.

There are some well-known improvements on the original SSA. For example, condition (C_1) can be relaxed to $D \mid 4n$ using the $\sqrt{2}$ trick; see [12].

3.1 Memory Requirement

For multiplying two N-bit integers, the total input size is $2N$ bits. The output size is $2N$ bits as well. The values of $\text{dft}(\mathbf{a})$ and $\text{dft}(\mathbf{b})$ in equation (2) are the intermediate results. We have

$$\text{dft}(\mathbf{a}) \in \left(\frac{\mathbb{Z}}{(2^n + 1)\mathbb{Z}} \right)^D,$$

a D-dimension tuple. By condition (C_2), the size is

$$|\text{dft}(\mathbf{a})| \ge nD \ge (2M + k)D > 4MK = 4N.$$

It is possible to choose the SSA parameter such that the DFT size is within the $(4N, 5N)$ interval. Therefore, the total space required for the multiplication is in $(12N, 14N]$ if the buffers for the input, the intermediate results and the output are separated. For an in-place implementation, i.e. the buffers are reused, the total required space decreases to $(8N, 10N]$. As an example, the required memory is from 1 TB to 1.25 TB for multiplying two 1-terabit integers with an in-place implementation; see also §4.4 and Table 2.

4. MapReduce-SSA

In our design, we target on input integers in terabit scale. SSA is a recursive algorithm. The first level call has to process multi-terabit data. The parameters are selected such that the data in the recursive calls is only in megabit or gigabit scale. Therefore, the recursive calls can be executed in a single machine locally. MapReduce framework is used to handle the first level call so that the computation is distributed across the machines and executed in parallel.

4.1 Data Model

In order to allow accessing terabit integers efficiently, an integer is divided into sequences of records. An integer x is represented as a D-dimensional tuple

$$\mathbf{x} = (x_{D-1}, x_{D-2}, \ldots, x_0) \in \mathbb{Z}^D$$

as illustrated in §3. It is customary to call the components of $\mathbf{x}$ *the digits* of x. Let

$$D = IJ. \tag{3}$$

where $I > 1$ and $J > 1$ are powers of two. For $0 \le i < I$, define a J-dimensional tuple

$$\mathbf{x}^{(i)} \stackrel{\text{def}}{=} (x_{(J-1)I+i}, x_{(J-2)I+i}, \ldots, x_i) \in \mathbb{Z}^J \tag{4}$$

so that

$$
\begin{pmatrix} \mathbf{x}^{(0)} \\ \mathbf{x}^{(1)} \\ \vdots \\ \mathbf{x}^{(I-1)} \end{pmatrix} = \begin{pmatrix} x_{(J-1)I} & x_{(J-2)I} & \cdots & x_0 \\ x_{(J-1)I+1} & x_{(J-2)I+1} & \cdots & x_1 \\ \vdots & \vdots & \ddots & \vdots \\ x_{(J-1)I+(I-1)} & x_{(J-2)I+(I-1)} & \cdots & x_{I-1} \end{pmatrix}.
$$

We call $(\mathbf{x}^{(0)}, \ldots, \mathbf{x}^{(I-1)})^t$ the (I, J)-format of $\mathbf{x}$.

Each $\mathbf{x}^{(i)}$ is represented as a sequence of J records and each record is a key-value pair, where the key is the index and the value is the value of the digit as shown in Table 1. There are I sequences for the entire tuple $\mathbf{x}$.

Record #	<Key, Value>
0	$<i, \quad x_i>$
1	$<J+i, \quad x_{J+i}>$
$\vdots$	$\vdots$
$J-1$	$<(J-1)I+i, \quad x_{(J-1)I+i}>$

Table 1: The representation of $\mathbf{x}^{(i)}$

We give a few definitions in the following. A tuple

$$\mathbf{x} = (x_{D-1}, \ldots, x_0)$$

is *normalized* if $0 \le x_t < 2^M$ for all $0 \le t < D$. The input tuples of the algorithm are assumed to be normalized and the output tuple returned is also normalized. Define the *left component-shift* operator $\ll$, for $m \ge 0$,

$$\mathbf{x} \ll m \stackrel{\text{def}}{=} (x_{D-m-1}, x_{D-m-2}, \ldots, x_0, \underbrace{0, \ldots, 0}_{m}). \quad (5)$$

By definition, $\mathbf{x} \ll 0$ is a no-op. Let $\mathbf{y} = (y_{D-1}, \ldots, y_0)$. Define the *addition with carrying* operator $\oplus$ as below,

$$\mathbf{x} \oplus \mathbf{y} \stackrel{\text{def}}{=} (s_{D-1} \bmod 2^M, \ldots, s_0 \bmod 2^M), \quad (6)$$

where $s_{-1} = 0$ and $s_t = x_t + y_t + \left\lfloor \frac{s_{t-1}}{2^M} \right\rfloor$ for $0 \le t < D$.

4.2 Algorithms

Using the notations in §3, SSA consists of four steps,

S1: two forward FFTs, $\hat{\mathbf{a}} \stackrel{\text{def}}{=} \text{dft}(\mathbf{a})$ and $\hat{\mathbf{b}} \stackrel{\text{def}}{=} \text{dft}(\mathbf{b})$;

S2: componentwise multiplication, $\hat{\mathbf{p}} \stackrel{\text{def}}{=} \hat{\mathbf{a}} * \hat{\mathbf{b}}$;

S3: a backward FFT, $\mathbf{p} = \text{dft}^{-1}(\hat{\mathbf{p}})$; and

S4: carrying, normalizing $\mathbf{p}$.

S1 first reads the two input integers, and then runs two forward FFTs in parallel. Forward FFTs are performed by the *MapReduce-FFT* algorithm described in §4.2.2. It is clear that the componentwise multiplication in S2 can be executed in parallel. Notice that the digits in $\hat{\mathbf{a}}$ and $\hat{\mathbf{b}}$ are relatively small in size. For terabit integers, the size of the digits is only in the scale of megabits. Thus, multiplications with these digits can be computed by a single machine. We discuss componentwise multiplication in §4.2.3. In S3, the

backward FFT, which is very similar to forward FFT, is again calculated by MapReduce-FFT. The carrying in S4 can be done by *MapReduce-Sum* as shown in §4.2.4.

4.2.1 Parallel-FFT

There is a well-known algorithm, called *parallel-FFT*, for evaluating a DFT in parallel. As before, let

$$\hat{\mathbf{a}} \stackrel{\text{def}}{=} (\hat{a}_{D-1}, \hat{a}_{D-2}, \ldots, \hat{a}_0) \stackrel{\text{def}}{=} \text{dft}(\mathbf{a}).$$

By the definition of DFT,

$$\hat{a}_j = \sum_{i=0}^{D-1} a_i \zeta^{ij} \qquad \text{for } 0 \le i < D, \quad (7)$$

where ζ is a principle Dth root of unity in $\mathbb{Z}/(2^n + 1)\mathbb{Z}$. We may take

$$\zeta = 2^{2n/D} \pmod{2^n + 1}.$$

Rewrite the summation, for all $0 \le j_0 < J$ and $0 \le j_1 < I$,

$$
\begin{aligned}
\hat{a}_{j_1 J + j_0} &= \sum_{i_0=0}^{I-1} \sum_{i_1=0}^{J-1} \zeta^{(i_1 I + i_0)(j_1 J + j_0)} a_{i_1 I + i_0} \\
&= \sum_{i_0=0}^{I-1} \zeta_I^{i_0 j_1} \left(\zeta^{i_0 j_0} \sum_{i_1=0}^{J-1} \zeta_J^{i_1 j_0} a_{i_1 I + i_0} \right),
\end{aligned}
$$

where $\zeta_I \stackrel{\text{def}}{=} \zeta^J$ and $\zeta_J \stackrel{\text{def}}{=} \zeta^I$ are principle Ith and Jth roots of unity in $\mathbb{Z}/(2^n + 1)\mathbb{Z}$, respectively. For $0 \le i < I$, let

$$\mathbf{a}^{(i)} \stackrel{\text{def}}{=} (a_{(J-1)I+i}, a_{(J-2)I+i}, \ldots, a_i)$$

and

$$\widehat{\mathbf{a}^{(i)}} \stackrel{\text{def}}{=} \left(\widehat{a^{(i)}}_{J-1}, \widehat{a^{(i)}}_{J-2}, \ldots, \widehat{a^{(i)}}_0 \right) \stackrel{\text{def}}{=} \text{dft}(\mathbf{a}^{(i)}). \quad (8)$$

These I DFTs with J points can be calculated in parallel. For $0 \le i < I$ and $0 \le j < J$, define

$$z_{jI+i} \stackrel{\text{def}}{=} \zeta^{ij} \widehat{a^{(i)}}_j, \quad (9)$$

$$\mathbf{z}^{[j]} \stackrel{\text{def}}{=} (z_{jI+(I-1)}, z_{jI+(I-2)}, \ldots, z_{jI}), \quad (10)$$

$$\widehat{\mathbf{z}^{[j]}} \stackrel{\text{def}}{=} \text{dft}(\mathbf{z}^{[j]}). \quad (11)$$

Notice that we use $[\,\cdot\,]$ in equation (10) in order to emphasize the difference between $\mathbf{z}^{[j]}$ and $\mathbf{z}^{(j)}$. The J DFTs with I points in equation (11) can be calculated in parallel. Finally,

$$\hat{a}_{j_1 J + j_0} = \sum_{i_0=0}^{I-1} \zeta_I^{i_0 j_1} z_{j_0 I + i_0} = \widehat{z^{[j_0]}}_{j_1}.$$

4.2.2 MapReduce-FFT

We may carry out parallel-FFT on MapReduce. For forward FFT, the input $\mathbf{a}$ is in (I, J)-format — it is divided into $\mathbf{a}^{(i)}$ for $0 \le i < I$ as in §4.1. There are I map and J reduce tasks in a job. The map tasks execute I parallel J-point FFTs as shown in equation (8), while the reduce tasks run J parallel I-point FFTs as shown in equation (11). For $0 \le j < J$, the output from reduce task j is

$$\hat{\mathbf{a}}^{(j)} = (\hat{a}_{(I-1)J+j}, \hat{a}_{(I-2)J+j}, \ldots, \hat{a}_j) = \widehat{\mathbf{z}^{[j]}}.$$

The output $\hat{\mathbf{a}}$ is written in (J, I)-format. The mapper and reducer algorithms are shown in Algorithms 1 and 2. Note that equation (9) computed in *(f.m.3)* indeed is a bit-shifting

since ζ is a power of two. Note also that the transposition of the elements from $\zeta^{ij}\widehat{\mathbf{a}^{(i)}}$ in equation (9) to $z^{[j]}$ in equation (10) is accomplished by the shuffle process, i.e. the transition from map tasks to reduce tasks.

ALGORITHM 1 (FORWARD FFT, MAPPER).

(f.m.1) read key i, value $\mathbf{a}^{(i)}$;

(f.m.2) calculate a J-point DFT by equation (8);

(f.m.3) componentwise bit-shift by equation (9);

(f.m.4) for $0 \leq j < J$, emit key j, value (i, z_{jI+i}).

ALGORITHM 2 (FORWARD FFT, REDUCER).

(f.r.1) receive key j, list $[(i, z_{jI+i})]_{0 \leq i < I}$;

(f.r.2) calculate an I-point DFT by equation (11);

(f.r.3) write key j, value $\widehat{\mathbf{z}^{[j]}}$.

Backward FFT is similar to forward FFT except that (1) it uses ζ^{-1} instead of ζ in the computation, and (2) it has an extra scale multiplication by $\frac{1}{D}$. In backward FFT, the roles of I and J are interchanged. The input is $\hat{\mathbf{p}}$ in (J, I)-format and the output $\mathbf{p}$ is in (I, J)-format, where $\mathbf{p} = \mathrm{dft}^{-1}(\hat{\mathbf{p}})$. There are J map tasks executing J parallel I-point FFTs and I reduce tasks running I parallel J-point FFTs. We skip the detailed equations and provide brief mapper and reducer algorithms in Algorithms 3 and 4. Similar to (f.m.3), we may perform (b.m.3) and (b.r.3) by bit-shifting since ζ^{-1} and D are also powers of two.

ALGORITHM 3 (BACKWARD FFT, MAPPER).

(b.m.1) read key j, value $\hat{\mathbf{p}}^{(j)}$;

(b.m.2) calculate an I-point DFT inverse;

(b.m.3) perform a componentwise bit-shifting;

(b.m.4) emit results.

ALGORITHM 4 (BACKWARD FFT, REDUCER).

(b.r.1) receive the corresponding map outputs;

(b.r.2) calculate a J-point DFT inverse;

(b.r.3) componentwise bit-shift for scaling by $\frac{1}{D}$;

(b.r.4) write key i, value $\mathbf{p}^{(i)}$.

We end the section by discussing the data size involved in MapReduce-FFT. Suppose the input is an integer with 1 TB (i.e. $N = 2^{43}$) and, say, $D = 2^{22}$. Then, the exponent in the Schönhage-Strassen modulus is

$$ n = \frac{5D}{2} = 10,485,760. $$

The size of each component in the tuple is 1.25 MB. If

$$ I = J = 2^{11} = 2048, $$

each task, map or reduce, processes a 2048-dimensional tuple with 2.5 GB data uniformly in all FFTs. If I and J are distinct, the machines running the tasks with smaller count have to process more data. Suppose $I = 4096$ and $J = 1024$. In the forward FFTs, each map task processes only 1.25 GB data but each reduce task has to process 5 GB data. On the other hand, each map and reduce task respectively processes 5 GB and 1.25 GB data in the backward FFT. See also Table 3 in §4.4.

4.2.3 Componentwise Multiplication

Componentwise multiplication can be easily done by a map-only job with J map tasks and zero reduce tasks. The mapper j reads $\hat{\mathbf{a}}^{(j)}$ and $\hat{\mathbf{b}}^{(j)}$, computes

$$ \hat{\mathbf{p}}^{(j)} \stackrel{\text{def}}{=} \hat{\mathbf{a}}^{(j)} * \hat{\mathbf{b}}^{(j)} \tag{12} $$

over $\mathbb{Z}/(2^n+1)\mathbb{Z}$ and then writes $\hat{\mathbf{p}}^{(j)}$. Componentwise multiplication indeed can be incorporated in Algorithm 3 to eliminate the extra map-only job. We replace Step (b.m.1) with the following.

ALGORITHM 5 (COMPONENTWISE-MULT).

(b.m.1.1) read key j, value $(\hat{\mathbf{a}}^{(j)}, \hat{\mathbf{b}}^{(j)})$;

(b.m.1.2) compute $\hat{\mathbf{p}}^{(j)}$ by equation (12).

Note that a careful implementation can avoid loading both $\hat{\mathbf{a}}^{(j)}$ and $\hat{\mathbf{b}}^{(j)}$ in the memory at the same time. We may read, multiply and deallocate the digits pair-by-pair.

4.2.4 Carrying

Since $\mathbf{p} = (p_{D-1}, \ldots, p_0)$ is the cyclic convolution of $\mathbf{a}$ and $\mathbf{b}$, we have

$$ 0 \leq p_t < 2^{2M+k} \qquad \text{for } 0 \leq t < D. $$

The tuple $\mathbf{p}$ is not normalized and the carries may affect one or more following components. The *carrying* operation is to normalize $\mathbf{p}$. It is straightforward to perform carrying sequentially but non-trivial to execute it in parallel efficiently. We reduce carrying to a summation of three normalized tuples. The summation can be done by *MapReduce-Sum* given in §4.2.5.

It is obvious that SSA is inefficient for $k \geq M$, so we only have to consider $k < M$. In practice, k is much smaller than M. For $0 \leq t < D$, write

$$ p_t = p_{t,2}2^{2M} + p_{t,1}2^M + p_{t,0} \tag{13} $$

such that $0 \leq p_{t,s} < 2^M$ for $s = 0, 1, 2$. Recall that $\ll$ and $\oplus$ are the left component-shift and the addition with carrying operators defined in equations (5) and (6). Let

$$ \mathbf{p}_{[s]} \stackrel{\text{def}}{=} (p_{D-1,s}, \ldots, p_{0,s}) \ll s \tag{14} $$

be D-dimensional normalized tuples. The dropped digits are

$$ p_{D-1,1} = p_{D-1,2} = p_{D-2,2} = 0. $$

If these digits are non-zero, the integer product $p > 2^{2N}$. The sum

$$ \tilde{\mathbf{p}} \stackrel{\text{def}}{=} \mathbf{p}_{[0]} \oplus \mathbf{p}_{[1]} \oplus \mathbf{p}_{[2]} \tag{15} $$

is the normalized form of $\mathbf{p}$. With equations (4), (13) and (14), it is not hard to see that the tuples

$$ \mathbf{p}_{[0]}^{(i)}, \quad \mathbf{p}_{[1]}^{((i+1) \bmod I)} \quad \text{and} \quad \mathbf{p}_{[2]}^{((i+2) \bmod I)} $$

can be constructed by rearranging the digits of $\mathbf{p}^{(i)}$. The steps for splitting $\mathbf{p}^{(i)}$ into three tuples can be incorporated in Algorithm 4. We replace Step (b.r.4) with the following.

ALGORITHM 6 (SPLITTING).

(b.r.4.1) construct the tuples $\mathbf{p}_{[0]}^{(i)}$, $\mathbf{p}_{[1]}^{((i+1) \bmod I)}$ and $\mathbf{p}_{[2]}^{((i+2) \bmod I)}$ from $\mathbf{p}^{(i)}$ by equations (4), (13) and (14);

(b.r.4.2) write key i, value $(\mathbf{p}_{[0]}^{(i)}, \mathbf{p}_{[1]}^{((i+1) \bmod I)}, \mathbf{p}_{[2]}^{((i+2) \bmod I)})$.

4.2.5 MapReduce-Sum

Let $m > 1$ be an integer constant, which is much smaller then N. For $0 \leq s < m$, let

$$\mathbf{y}_{[s]} \stackrel{\text{def}}{=} (y_{D-1,s}, \ldots, y_{0,s})$$

be m normalized tuples. Extending the definition (6) of addition with carrying, define *sum with carrying*

$$\mathbf{y}_{[0]} \oplus \cdots \oplus \mathbf{y}_{[m-1]} \stackrel{\text{def}}{=} (s_{D-1} \bmod 2^M, \ldots, s_0 \bmod 2^M),$$

where $s_{-1} = 0$ and, for $0 \leq t < D$,

$$s_t = y_{t,0} + \cdots + y_{t,m-1} + \left\lfloor \frac{s_{t-1}}{2^M} \right\rfloor. \tag{16}$$

If $\mathbf{y}_{[s]}$'s are considered as integers $y_{[s]}$ as in §3, then $\mathbf{y}_{[0]} \oplus \cdots \oplus \mathbf{y}_{[m-1]}$ is the normalized representation of the integer sum

$$\sum_{s=0}^{m-1} y_{[s]} \pmod{2^{DM}}.$$

The classical addition algorithm is a sequential algorithm because the carries in the higher digits depend on the carries in the lower digits. It is not suitable to our environment since the entire D-dimensional tuple is in terabyte scale. Processing it with a single machine takes too much time.

An observation is that it is unnecessary to have the entire digit sequence for determining the carries $\left\lfloor \frac{s_{t-1}}{2^M} \right\rfloor$ in equation (16). The carries can be obtained using a much smaller data set. For $0 \leq t < D$, let

$$z_t \stackrel{\text{def}}{=} y_{t,0} + \cdots + y_{t,m-1}, \tag{17}$$
$$r_t \stackrel{\text{def}}{=} z_t \bmod 2^M, \tag{18}$$
$$c_t \stackrel{\text{def}}{=} \left\lfloor \frac{z_t}{2^M} \right\rfloor \tag{19}$$

be the componentwise sums, the remainders and the intermediate carries, respectively. Note that c_t may not equals to $\left\lfloor \frac{s_t}{2^M} \right\rfloor$ since, when c_t's are added to the digits, overflow may occur. Then carrying is required again in such case. Define the *difference*

$$d_t = \begin{cases} 2^M - r_t, & \text{if } 2^M - r_t < m; \\ m, & \text{otherwise.} \end{cases} \tag{20}$$

Notice that $0 < d_t \leq m$. We also have

$$0 \leq c_t \leq \left\lfloor \frac{s_t}{2^M} \right\rfloor < m. \tag{21}$$

The last inequality is due to the fact that the input tuples $\mathbf{y}_{[s]}$'s are normalized. The sizes of the tuples $(c_{D-1}, \cdots, c_0)$ and $(d_{D-1}, \cdots, d_0)$ are in $O(D)$, given m a constant by assumption. Similar to [12, equation (3)], we have

$$D \leq \sqrt{8N}. \tag{22}$$

The data size is reduced from $O(N)$, for classical addition algorithm, to $O(D) = O(\sqrt{N})$; i.e. from terabits to megabits. So these two tuples can be sequentially processed in a single machine efficiently. For $0 \leq t < D$, let

$$c'_t = \begin{cases} c_t + 1, & \text{if } t > 0 \text{ and } c'_{t-1} \geq d_t; \\ c_t, & \text{otherwise.} \end{cases} \tag{23}$$

Then, $c'_0, \ldots, c'_{D-1}$ are the carries as shown below.

THEOREM 7. *Let*

$$y_t = \begin{cases} r_0, & \text{if } t = 0; \\ (r_t + c'_{t-1}) \bmod 2^M, & \text{if } 1 \leq t < D. \end{cases} \tag{24}$$

Then,

$$(y_{D-1}, \ldots, y_0) = \mathbf{y}_{[0]} \oplus \cdots \oplus \mathbf{y}_{[m-1]}.$$

PROOF. We prove by induction that $c'_t = \left\lfloor \frac{s_t}{2^M} \right\rfloor$ for $0 \leq t < D$. Clearly,

$$c'_0 = c_0 = \left\lfloor \frac{z_0}{2^M} \right\rfloor = \left\lfloor \frac{s_0}{2^M} \right\rfloor.$$

Assume $c'_t = \left\lfloor \frac{s_t}{2^M} \right\rfloor$ for $0 \leq t < D - 1$. By equation (21), $c'_t < m$. Then,

$$s_{t+1} = z_{t+1} + \left\lfloor \frac{s_t}{2^M} \right\rfloor = c_{t+1}2^M + r_{t+1} + c'_t \tag{25}$$

by equations (16), (18) and (19). Since m is much smaller than 2^M by assumption, we have

$$0 \leq r_{t+1} + c'_t < 2^M + m < 2^{M+1}.$$

Then,

$$\left\lfloor \frac{r_{t+1} + c'_t}{2^M} \right\rfloor = \begin{cases} 1, & \text{if } r_{t+1} + c'_t \geq 2^M; \\ 0, & \text{otherwise.} \end{cases} \tag{26}$$

Suppose $d_{t+1} = m$. By equation (20), $2^M - r_{t+1} \geq m$. We have

$$r_{t+1} + c'_t = r_{t+1} + \left\lfloor \frac{s_t}{2^M} \right\rfloor < r_{t+1} + m \leq 2^M.$$

Divide equation (25) by 2^M,

$$\left\lfloor \frac{s_{t+1}}{2^M} \right\rfloor = c_{t+1} + \left\lfloor \frac{r_{t+1} + c'_t}{2^M} \right\rfloor = c_{t+1} = c'_{t+1}$$

by equations (23) and (26). Suppose $d_{t+1} < m$. Then, equation (23) becomes

$$c'_{t+1} = \begin{cases} c_{t+1} + 1, & \text{if } r_{t+1} + c'_t \geq 2^M; \\ c_{t+1}, & \text{otherwise;} \end{cases}$$

by substituting $d_{t+1} = 2^M - r_{t+1}$. By equations (26) and (25),

$$c'_{t+1} = c_{t+1} + \left\lfloor \frac{r_{t+1} + c'_t}{2^M} \right\rfloor = \left\lfloor \frac{s_{t+1}}{2^M} \right\rfloor.$$

Finally, by equation (24),

$$y_0 = r_0 \equiv z_0 \equiv s_0 \pmod{2^M}$$

and, for $1 \leq t < D$,

$$y_t \equiv r_t + c'_{t-1} \equiv z_t + \left\lfloor \frac{s_{t-1}}{2^M} \right\rfloor \equiv s_t \pmod{2^M}.$$

The theorem follows. $\square$

We present MapReduce-Sum below. All tuples are represented in the (I, J)-format described in §4.1. Two jobs are involved in MapReduce-Sum. The first job for componentwise summation (Algorithm 8) is a map-only job, which has I map and zero reduce tasks. The mapper i reads input tuples $\mathbf{y}_{[s]}^{(i)}$ for $0 \leq s < m$, and writes the remainders $\mathbf{r}^{(i)}$, the intermediate carries $\mathbf{c}^{(i)}$ and the differences $\mathbf{d}^{(i)}$.

ALGORITHM 8 (COMPONENTWISE SUM, MAPPER).

(s.m.1) *read key i, value $(\mathbf{y}_{[0]}{}^{(i)}, \ldots, \mathbf{y}_{[m]}{}^{(i)})$;*

(s.m.2) *compute $\mathbf{z}^{(i)}$ by equation (17);*

(s.m.3) *compute $\mathbf{r}^{(i)}$ by equation (18);*

(s.m.4) *compute $\mathbf{c}^{(i)}$ by equation (19);*

(s.m.5) *compute $\mathbf{d}^{(i)}$ by equation (20);*

(s.m.6) *write key i, value $(\mathbf{r}^{(i)}, \mathbf{c}^{(i)}, \mathbf{d}^{(i)})$.*

The second job for carrying (Algorithms 9 and 10) has a single map task and I reduce tasks. The mapper reads all the intermediate carries and all the differences, and then evaluates the actual carries. The reducers receive the actual carries from the mapper, read the corresponding remainders and then compute the final output. For single-map jobs, we may use the *null key*, denoted by $\emptyset$, as the input key.

ALGORITHM 9 (CARRYING, MAPPER).

(c.m.1) *read key $\emptyset$, value $\left\{ (i, \mathbf{c}^{(i)}, \mathbf{d}^{(i)}) : 0 \leq i < I \right\}$;*

(c.m.2) *compute $\mathbf{c}'$ by equation (23);*

(c.m.3) *for $0 \leq i < I$, emit key i, value $\mathbf{c}'^{((i-1) \bmod I)}$.*

ALGORITHM 10 (CARRYING, REDUCER).

(c.r.1) *receive key i, singleton list $[\mathbf{c}'^{((i-1) \bmod I)}]$;*

(c.r.2) *read $\mathbf{r}^{(i)}$;*

(c.r.3) *compute $\mathbf{y}^{(i)}$ by equation (24);*

(c.r.4) *write key i, value $\mathbf{y}^{(i)}$.*

4.3 Summary

The input of MapReduce-SSA is two integers a and b, which are represented as normalized tuples $\mathbf{a}$ and $\mathbf{b}$ in (I, J)-format. The output is also in (I, J)-format, a normalized tuple $\tilde{\mathbf{p}}$, which represents an integer p such that $p = ab$. MapReduce-SSA has the following sequence of jobs:

J1: two concurrent forward FFT jobs (Alg. 1 & 2);

J2: a backward FFT job with componentwise multiplication and splitting (Alg. 3 with 5, & 4 with 6);

J3: a componentwise summation job (Alg. 8);

J4: a carrying job (Alg. 9 & 10).

4.4 Choosing Parameters

In SSA, once D, the dimension of the FFTs, is fixed, all other variables, including the exponent n in the Schönhage-Strassen modulus $2^n + 1$, are fixed consequently. The standard choice is $D \approx \sqrt{N}$.

In MapReduce-SSA, there is an additional parameter I, defined in equation (3). It determines the number of sequences in the integer representation, the numbers of tasks in the jobs and the memory required in each task. Therefore, the choice of I depends on the cluster capacity, i.e. the numbers of map and reduce tasks supported, and the available memory in each machine.

Denote the number of map and reduce tasks in the forward FFTs by $N_{f,m}$ and $N_{f,r}$, and the number of map and reduce tasks in the backward FFT by $N_{b,m}$ and $N_{b,r}$. Then,

$$N_{f,m} = N_{b,r} = I \qquad \text{and} \qquad N_{f,r} = N_{b,m} = J.$$

Excluding the memory overhead from the system software and the temporary variables, let $M_{f,m}$ and $M_{f,r}$ be the number of bits of the required memory in map and reduce tasks in forward FFT, and $M_{b,m}$ and $M_{b,r}$ be the number of bits of the required memory in map and reduce tasks in backward FFT. Then,

$$M_{f,m} = M_{b,r} = Jn \qquad \text{and} \qquad M_{f,r} = M_{b,m} = In.$$

We always choose $I \leq J$ so that $M_{b,m} = In \leq Jn$, in order to have more memory available for running Algorithm 5 in the map tasks of the backward FFT. In Tables 2, 3 and 4, we show some possible choices of the parameters D and I, and the corresponding values of some other variables for $N = 2^{40}$, 2^{43} and 2^{46}. The memory requirement is just 0.56 GB per task for $N = 2^{40}$. For $N = 2^{46}$, it requires 12 GB per task.

$N = 2^{40}$	J	n	In	Jn
$D = 2^{19}, I = 2^9$	2^{10}	$16.5D$	0.52GB	1.03GB
$D = 2^{20}, I = 2^{10}$	2^{10}	$4.5D$	0.56GB	0.56GB
$D = 2^{21}, I = 2^{10}$	2^{11}	$1.5D$	0.38GB	0.75GB
$D = 2^{22}, I = 2^{11}$	2^{11}	$0.5D$	0.50GB	0.50GB

Table 2: Possible parameters for $N = 2^{40}$

$N = 2^{43}$	J	n	In	Jn
$D = 2^{20}, I = 2^{10}$	2^{10}	$32.5D$	4.06GB	4.06GB
$D = 2^{21}, I = 2^{10}$	2^{11}	$8.5D$	2.13GB	4.25GB
$D = 2^{22}, I = 2^{11}$	2^{11}	$2.5D$	2.50GB	2.50GB
$D = 2^{23}, I = 2^{11}$	2^{12}	D	2.00GB	4.00GB

Table 3: Possible parameters for $N = 2^{43}$

$N = 2^{46}$	J	n	In	Jn
$D = 2^{22}, I = 2^{11}$	2^{11}	$16.5D$	16.5GB	16.5GB
$D = 2^{23}, I = 2^{11}$	2^{12}	$4.5D$	9.0GB	18.0GB
$D = 2^{24}, I = 2^{12}$	2^{12}	$1.5D$	12.0GB	12.0GB
$D = 2^{25}, I = 2^{12}$	2^{13}	$0.5D$	8.0GB	16.0GB

Table 4: Possible parameters for $N = 2^{46}$

4.5 A Prototype Implementation

The program *DistMpMult* is a prototype implementation, which includes *DistFft*, *DistCompSum* and *DistCarrying* as

subroutines. DistFft implements Algorithms 1, 2, 3 and 4 with Algorithms 5 and 6, while DistCompSum implements Algorithm 8 and DistCarrying implements Algorithms 9 and 10. We have verified the correctness of the implementation by comparing the outputs from DistMpMult to the outputs of the GMP library.

Since DistFft is notably more complicated than DistComp-Sum and DistCarrying, we have measured the performance of DistFft on a 1350-node cluster. Each node has 6 GB memory, and supports 2 map tasks and 1 reduce task. The cluster is installed with Apache Hadoop version 0.20. It is a shared cluster, which is used by various research projects in Yahoo!. Figure 3 shows a graph of the elapsed time t against the input integer size N in base-2 logarithmic scales. The elapsed time includes the computation of two parallel forward FFTs and then a backward FFT. From the range $2^{32} \leq N \leq 2^{36}$, the elapsed time fluctuates between 2 to 4 minutes because of the system overheads. The utilization is extremely low in these cases. From the range $2^{36} \leq N \leq 2^{40}$, the graph is almost linear with slope $\approx 0.8 < 1$, which indicates that the cluster is still underutilized since FFT is an essentially linear time algorithm. Unfortunately, the 1350-node cluster is unable to run DistFft for $N \geq 2^{41}$ due to the memory limitation configured in the cluster, so that it limits the aggregated memory usage of individual jobs.
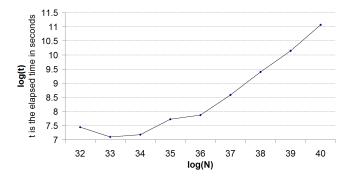

Figure 3: Actual running time of DistFft on a 1350-node cluster

For $N = 2^{40}$, we choose $D = 2^{20}$ and $I = 2^{10}$. Then, $J = I$ and

$$n = 4.5D = 4,718,592$$

as shown in Table 2. In a forward FFT job, the average durations of a map, a shuffle and a reduce task are 157, 223 and 90 seconds, respectively. Although the total is only 7.83 minutes, the the job duration is 15.38 minutes due to scheduling and other overheads. In a backward FFT job, the average durations for a map, a shuffle and a reduce task are 251, 483 and 77 seconds, and the job duration is 18.17 minutes. The itemized details are shown in Table 5. It shows that the computation is I/O bound. The CPU intensive computation only contributes 10% to 15% of the elapsed time excluding the system overheads.

Since there are already benchmarks showing that Hadoop is able to scale adequately on manipulating 100-TB data on a large cluster [14], we estimate that DistFft will be able to process integers with $N = 2^{46}$ in some suitable clusters. Such integers have 21 trillion decimal digits and require 8 TB space. A suitable cluster may consists of 4200 nodes with

$N = 2^{40}$		Step	Duration
Forward FFT	Mapper (Alg. 1)	(f.m.1)	12s
		(f.m.2)	24s
		(f.m.3)	1s
		(f.m.4)	120s
		Total	157s
	Shuffle		223s
	Reducer (Alg. 2)	(f.r.1)	19s
		(f.r.2)	22s
		(f.r.3)	49s
		Total	90s
Backward FFT	Mapper (Alg. 3)	(b.m.1)	89s
		(b.m.2)	70s
		(b.m.3)	2s
		(b.m.4)	90s
		Total	251s
	Shuffle		483s
	Reducer (Alg. 4)	(b.r.1)	20s
		(b.r.2)	20s
		(b.r.3)	24s
		(b.r.4)	13s
		Total	77s

Table 5: Average duration for each step in MapReduce-FFT with $N = 2^{40}$

16 GB memory in each machine or 2100 nodes with 32 GB memory in each machine according to the parameters suggested in Table 4.

5. CONCLUSION

An integer multiplication algorithm, *MapReduce-SSA*, is designed for multiplying integers in terabit scale on clusters with many commodity machines. It employs the ideas from Schönhage-Strassen algorithm in the MapReduce model. We first introduce a data model which allows efficiently accessing terabit integers in parallel. We show *forward/backward MapReduce-FFT* for calculating DFT and the inverse DFT, and describe how to perform componentwise multiplications and carrying operations. Componentwise multiplication is simply incorporated in the backward FFT step. The carrying operation is reduced to a summation, which is computed by *MapReduce-Sum*. A MapReduce-SSA computation involves five jobs — two concurrent jobs for two forward FFTs, a job for a backward FFT, and two jobs for carrying. In addition, we discuss the choices of the parameters and the impacts on cluster capacity and memory requirements.

We have developed a program, *DistMpMult*, which is a prototype implementation of the entire MapReduce-SSA algorithm. It includes *DistFft*, *DistCompSum* and *DistCarry-*

ing as subroutines for calculating DFT/inverse DFT, componentwise summation and carrying, respectively. All routines run on Apache Hadoop clusters. The correctness of the implementation has been verified by comparing the outputs from DistMpMult to the outputs of the GMP library. Our experiments demonstrate that DistFft is able to process terabit integers efficiently using a 1350-node cluster with 6 GB memory in each machine. With the supporting data from benchmarks on Hadoop, we estimate that DistFft is able to process integers with 8 TB in size, or, equivalently, 21 trillion decimal digits, on some large clusters.

6. REFERENCES

[1] F. Bellard. Pi computation record, 2009.
`http://bellard.org/pi/pi2700e9/announce.html`.

[2] F. Bellard. Tachuspi, 2009.
`http://bellard.org/pi/pi2700e9/tpi.html`.

[3] D. J. Bernstein. Multidigit multiplication for mathematicians, 2001. Available at
`http://cr.yp.to/papers/m3.pdf`.

[4] D. J. Bernstein. Fast multiplication and its applications. Number 44 in Mathematical Sciences Research Institute Publications, pages 325–384. Cambridge University Press, 2008.

[5] 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. 96(3):201–219, 1989.

[6] R. P. Brent and P. Zimmermann. *Modern Computer Arithmetic*. Number 18 in Cambridge Monographs on Computational and Applied Mathematics. Cambridge University Press, Cambridge, United Kingdom, 2010.

[7] D. Chudnovsky and G. Chudnovsky. The computation of classical constants. In *Proc. Nat. Acad. Sci. USA*, volume 86, pages 8178–8182. Academic Press, 1989.

[8] R. Crandall and C. Pomerance. *Prime Numbers: A Computational Perspective*. Springer-Verlag, New York, NY, 2001.

[9] A. De, P. P. Kurur, C. Saha, and R. Saptharishi. Fast integer multiplication using modular arithmetic. In *Proceedings of the 40th annual ACM Symposium on Theory of Computing*, pages 499–506. ACM, 2008.

[10] J. Dean and S. Ghemawat. MapReduce: simplified data processing on large clusters. In *Proceedings of the 6th conference on Symposium on Opearting Systems Design & Implementation*, pages 137–150. USENIX Association, 2004.

[11] M. Fürer. Faster integer multiplication. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 57–66. ACM, 2007.

[12] P. Gaudry, A. Kruppa, and P. Zimmermann. A GMP-based implementation of Schönhage-Strassen's large integer multiplication algorithm. In *Proceedings of the 2007 International Symposium on Symbolic and Algebraic Computation*, pages 167–174. ACM, 2007.

[13] D. E. Knuth. *The Art of Computer Programming, Volume 2: Seminumerical Algorithms*. Addison-Wesley, Reading, MA, 1997.

[14] O. O'Malley and A. C. Murthy. Winning a 60 second dash with a yellow elephant, 2009. Available at
`http://sortbenchmark.org/Yahoo2009.pdf`.

[15] A. Schönhage and V. Strassen. Schnelle Multiplikation großer Zahlen. *Computing*, 7:281–292, 1971.

[16] K. Shvachko, H. Kuang, S. Radia, and R. Chansler. The Hadoop Distributed File System. In *26th IEEE Symposium on Massive Storage Systems and Technologies*, 2010.

[17] D. Takahashi, 2010. Personal communication.

[18] D. Takahashi. Parallel implementation of multiple-precision arithmetic and 2,576,980,370,000 decimal digits of π calculation. *Parallel Comput.*, 36:439–448, August 2010.

[19] A. J. Yee, 2010. Personal communication.

[20] A. J. Yee. y-cruncher – a multi-threaded pi-program, 2010. `http://www.numberworld.org/y-cruncher/`.

[21] A. J. Yee and S. Kondo. 5 trillion digits of pi - new world record, 2010. `http://www.numberworld.org/misc_runs/pi-5t/announce_en.html`.