A Comparison of Two Methods for Montgomery Multiplication Using C++

Oskar Asplin asplin@umail.ucsb.edu

June 2017

Abstract

This paper will examine how to implement the classic Montgomery algorithm and the CIOS algorithm, and compare their runtime. Both algorithms will be implemented in the C++ programming language and tested using various sets of input parameters. The CIOS algorithm is stated to be more effective, so I will find out if this is true for my implementation, and to what degree they differ. Their runtime for unique sets of input will be compared and analyzed.

1 Introduction

Montgomery multiplication is a method for performing fast modular multiplication, and it was introduced by Peter L. Montgomery in 1985. It has made a huge difference in the world of binary computation and cryptographic engineering, by improving the efficiency of modular multiplication. The CIOS (Coarsely Integrated Operand Scanning) algorithm is another method of performing Montgomery multiplication, which is even more efficient. CIOS only requires 1.5 integer multiplications instead of 3 as the classic method needs.

In this project I have implemented the two algorithms in C++ in an attempt to further understand how they work, and test how they compare to each other in terms of runtime.

2 Implementation in C++

Before writing code, the two methods was first studied by reviewing professor Koç's course material in CS 293G, as well as Wikipedia and other academic sources online. This is the Pseudo code for the classic Montgomery multiplication:

function MonPro(a, b)input: a, b, n, r, n'output: $u = a \cdot b \cdot r^{-1} \mod n$ 1: $t \leftarrow a \cdot b$ 2: $m \leftarrow t \cdot n' \pmod{r}$ 3: $u \leftarrow (t + m \cdot n)/r$ 4: if $u \ge n$ then $u \leftarrow u - n$ 5: return u

It was fairly simple to code in C++, however I met several errors as I ran it with test input. By adjusting the code in order to fix the errors, I was able to make a solid code that can take any set of input parameters (a,b,n) of type int16 which returned the correct output u.

The Pseudo code for the CIOS method is as follows:

```
function MonPro(a, b)

input: a, b, n, r, n'

output: u = a \cdot b \cdot r^{-1} \mod n

1: u \leftarrow 0

2: for i = 0 to k - 1

3a: u \leftarrow u + A_i \cdot b

3b: u \leftarrow u + U_0 \cdot n

3c: u \leftarrow u/2

4: if u \ge n then u \leftarrow u - n

5: return u
```

The CIOS method was a bit easier to implement, but I could not resolve the problem of going bit by bit through A without doing a new byte multiplication. In hardware implementation of CIOS it is probably more effective than the classical method, but in my case, step 3 lead to more byte multiplications.

Both algorithms used the Extended Euclidean Algorithm which I made using a template I found online. Its Pseudo code is shown below:

$$(u, v) \leftarrow \text{EEA}(a, n)$$

 $u \cdot a - v \cdot n = 1$
 $a^{-1} = u \pmod{n}$

The code for the two methods is added in section 5.

3 Results

After successfully implementing both Montgomery methods and testing them with numerous inputs to see they worked properly, I made a function to time them. 100 000 random input sets were created and then fed into the two Montgomery functions. Both algorithms were timed separately, and the test was performed 200 times.

The results of randomly generated int16-inputs are shown in figure 1 and 2. As seen in the figures, the classic method performed much better than CIOS, which was approximately 71% slower. The median time for the classic Montgomery was 314 μ s, while the median for the CIOS was 539 μ s. This means that a single product calculation takes about 314 ns and 539 ns for the two methods. In theory the CIOS should perform better at a hardware level as it needs fewer multiplications, but due to my implementation in a higher level language (C++), the CIOS algorithm was not able to perform as fast as it could. I believe the multiplications at step 3 in the algorithm were performed as byte multiplications instead of going through A bitwise. I tried different ways of performing the third step, but was not able to improve the runtime.



Figure 1: Plot showing result of 200 tests with 100 000 random inputs



Figure 2: Boxplot showing result of 200 tests with 100 000 random inputs

When bounding the input to numbers up to 256 (int8), both methods performed better, and the CIOS method improved significantly. The results can be seen in figure 3 and 4. The classic method now had a median of 265 μ s and CIOS had a median of 333 μ s. Though CIOS is still 26% slower, its performance improved drastically. Again I believe this is due to step three in the algorithm, which now has to perform fewer multiplications. The number of operations for the classic method is constant in both cases.



Figure 3: Plot showing result of 200 tests with 100 000 random int8-inputs



Figure 4: Boxplot showing result of 200 tests with 100 000 random int8-inputs

All testing was done on a computer with an Intel Core i7-4500U processor with 1.8-3GHz performance.

4 Conclusion

This project has taught me how two different methods of performing Montgomery multiplication work in theory, and how they can be implemented using C++. As stated in the results, the CIOS is more effective if properly implemented in a lower programming language or at hardware level, however I was not able to achieve this ideal use of the algorithm. Therefore we can see significantly better results for the classic method in my C++ coding. Since I have little knowledge of how the compiler breaks down my written code, I do not know the exact number of multiplications or the exact execution in the compiled program. If one were to perfectly compare the difference between the two algorithms, one would have to code in an assembly language or know how the compiler operates.

5 C++ code

5.1 monPro.cpp

```
1 #include <iostream>
<sup>2</sup> using namespace std;
3 #include <bitset>
4 #include <stdio.h>
5 #include <math.h>
                           /* pow */
                           /* srand, rand */
6 #include <stdlib.h>
7 #include <sys/time.h>
                           /* gettimeofday, timeval */
8 #include <algorithm>
                           /* sort, begin, end */
9
10 #define NUM_INPUTS 100000
11 #define NUM_SAMPLES 200
12 \# define INT16\_SIZE 65536
13
14 // ax + by = gcd(a, b)
15 int gcdExtended(int a, int b, int *x, int *y)
16
  {
      // Base Case
17
      if (a == 0)
18
      {
19
           *x = 0;
20
          *y = 1;
21
          return b;
      }
23
24
      int x1, y1; // To store results of recursive call
25
      int gcd = gcdExtended(b\%a, a, \&x1, \&y1);
26
27
      // Update x and y using results of recursive call
28
      *x = y1 - (b/a) * x1;
      *y = x1;
30
31
32
      return gcd;
33
34
  // Find position of the MSB in Int16
35
  int msbPosInt(int a)
36
37
  ł
      38
      int pos = 16;
39
      int b = a \& bitCompare;
40
      while (b == 0)
41
42
      {
          bitCompare = bitCompare >> 1;
43
          pos -= 1;
44
          b = a \& bitCompare;
45
      }
46
      return pos;
47
48
49
```

```
50 int cios(int a, int b, int n)
51
   ł
       int k = msbPosInt(n)+1;
52
       int r = pow(2, k);
53
       bitset < sizeof(int) * 8 > a_bits(a);
54
       long u = 0;
56
       for (int i = 0; i < k; i++)
57
58
       ł
            // Step 3a
59
            u = u + a_bits[i]*b;
60
            // Step 3b
61
            u = u + (u \& 1) * n;
62
            // Step 3c
            u = u >> 1;
64
       }
65
       // Step 4
66
       if (u \ge n)
67
       {
68
            u = u - n;
       }
70
       return u;
71
   }
72
73
74 int classicMonPro(int a, int b, int n)
   {
75
       int k = msbPosInt(n)+1;
76
       int r = pow(2, k);
77
       int n_dot, x;
78
       unsigned int t;
79
       long u, m;
80
81
       gcdExtended(r, n, \&x, \&n_dot);
82
       if (x < 0) {
83
            x = n + x;
84
            n_dot = n_dot - r;
85
       }
86
       n_dot = -n_dot;
87
88
       // Step 1
89
       t = a * b;
90
91
       // Step 2
92
       m = t * n_dot;
93
       m = m \& (r-1);
94
95
       // Step 3
96
       u = (t + m*n);
97
       int msbR = msbPosInt(r); // divide by r
98
       u = u >> msbR;
99
100
       // Step 4
101
102
       if (u \ge n)
           u = u - n;
103
```

```
104
105
       // Step 5
106
107
       return u;
108
109
110 int main()
111
   {
112
       int u, v;
       struct timeval tvalBefore, tvalAfter;
       long classicTime [NUM_SAMPLES], ciosTime [NUM_SAMPLES];
114
       /* initialize random seed: */
       srand (time(NULL));
118
       int * n = new int [NUM_INPUTS];
119
       int * a = new int [NUM_INPUTS];
120
       int * b = new int [NUM_INPUTS];
       // Begin testing
       /****
124
                                         *****
       for (int test_x = 0; test_x < NUM_SAMPLES; test_x++)
       ł
126
            for (int i = 0; i < NUM_{INPUTS}; i++)
127
128
            ł
                n[i] = (rand() \% (INT16\_SIZE-1)/2)*2 + 1; // Need an odd
129
      number for gcd(r, n) = 1
                a[i] = rand() \% n[i];
130
                b[i] = rand() \% n[i];
            }
           // Record time for NUM_INPUTS of the classic montgomery
134
      algorithm
           gettimeofday (&tvalBefore, NULL);
136
137
            for (int i = 0; i < NUMINPUTS; i++)
            {
138
                u = classicMonPro(a[i], b[i], n[i]);
139
            }
140
141
            gettimeofday (&tvalAfter, NULL);
142
143
144
            classicTime[test_x] = ((tvalAfter.tv_sec - tvalBefore.tv_sec))
145
      *100000L
                   +tvalAfter.tv_usec) - tvalBefore.tv_usec;
146
            // Divide by a scalar to get average time per 1000:
147
            classicTime [test_x] /= 100L;
148
149
            // Record time for NUM_INPUTS of the CIOS algorithm
            gettimeofday (&tvalBefore, NULL);
152
            for (int i = 0; i < NUMINPUTS; i++)
153
154
```

```
\mathbf{v} = \operatorname{cios}(\mathbf{a}[\mathbf{i}], \mathbf{b}[\mathbf{i}], \mathbf{n}[\mathbf{i}]);
            }
156
157
             gettimeofday (&tvalAfter, NULL);
158
159
            ciosTime[test_x] = ((tvalAfter.tv_sec - tvalBefore.tv_sec)
160
       *100000L
                     +tvalAfter.tv_usec) - tvalBefore.tv_usec;
161
             // Divide by a ascalar to get average time per 1000:
162
            ciosTime [test_x] \neq 100L;
163
        }
164
        /**
                                                                               ******
165
        //End of tetsing
166
167
        /* Sort arrays */
168
        sort(classicTime , classicTime + NUM_SAMPLES);
        sort (ciosTime, ciosTime + NUM.SAMPLES);
170
171
        /* Print results */
172
        cout << "Time in microseconds for the two algorithms per 1000 inputs
173
       :" << endl;
        printf("classicMonPro min: %ld microseconds\nclassicMonPro max: %ld
174
       microsecondsn,
                  classicTime [0], classicTime [199]);
        printf("CIOS min: %ld microseconds\nCIOS max: %ld microseconds\n",
176
       ciosTime[0], ciosTime[199]);
        /*
178
        // Print all times
179
        cout << "Classic Montgomery times:" << endl;</pre>
180
        for (int \ i = 0; \ i < NUM_SAMPLES; \ i++)
181
        ł
182
            cout << classicTime[i] << endl;</pre>
183
184
185
        cout << "CIOS times:" << endl;</pre>
186
        for (int i = 0; i < NUM_SAMPLES; i++)
187
        ł
188
             cout << ciosTime[i] << endl;</pre>
189
190
        }
        */
191
192
        /* Delete dynamic memory usage */
193
        delete [] n;
194
        delete [] a;
195
        delete [] b;
196
197
        return 0;
198
199
   ł
```

Code/monPro.cpp