The Poly1305-AES message-authentication code

Daniel J. Bernstein *

Department of Mathematics, Statistics, and Computer Science (M/C 249)
The University of Illinois at Chicago
Chicago, IL 60607–7045
djb@cr.yp.to

Abstract. Poly1305-AES is a state-of-the-art message-authentication code suitable for a wide variety of applications. Poly1305-AES computes a 16-byte authenticator of a variable-length message, using a 16-byte AES key, a 16-byte additional key, and a 16-byte nonce. The security of Poly1305-AES is very close to the security of AES; the security gap is at most $14D\lceil L/16\rceil/2^{106}$ if messages have at most L bytes, the attacker sees at most 2^{64} authenticated messages, and the attacker attempts D forgeries. Poly1305-AES can be computed at extremely high speed: for example, fewer than $3.625(\ell + 170)$ Athlon cycles for an ℓ-byte message. This speed is achieved *without* precomputation; consequently, 1000 keys can be handled simultaneously without cache misses. Special-purpose hardware can compute Poly1305-AES at even higher speed. Poly1305-AES is parallelizable, incremental, and not subject to any intellectual-property claims.

1 Introduction

This paper introduces and analyzes Poly1305-AES, a state-of-the-art secret-key message-authentication code suitable for a wide variety of applications.

Poly1305-AES computes a 16-byte authenticator $\text{Poly1305}_r(m, \text{AES}_k(n))$ of a variable-length message m, using a 16-byte AES key k, a 16-byte additional key r, and a 16-byte nonce n. Section 2 of this paper presents the complete definition of Poly1305-AES.

Poly1305-AES has several useful features:

- **Guaranteed security if AES is secure.** The security gap is small, even for long-term keys. Assume, for example, that messages are packets up to 1024

* The author was supported by the National Science Foundation under grant CCR–9983950, and by the Alfred P. Sloan Foundation. Date of this document: 2004.11.01. Permanent ID of this document: 0018d9551b5546d97c340e0dd8cb5750. This is a preliminary version meant to announce ideas; it will be replaced by a final version meant to record the ideas for posterity. There may be big changes before the final version. Future readers should not be forced to look at preliminary versions, unless they want to check historical credits; if you cite a preliminary version, please repeat all ideas that you are using from it, so that the reader can skip it.

bytes; that the attacker sees 2^{64} messages authenticated under a Poly1305-AES key; that the attacker attempts a whopping 2^{75} forgeries; and that the attacker cannot break AES with probability above δ. Then, with probability at least $0.999999 - \delta$, *all* of the 2^{75} forgeries are rejected.

- **Cipher replaceability.** If anything does go wrong with AES, users can switch from Poly1305-AES to Poly1305-AnotherFunction, with an identical security guarantee. All the effort invested in the non-AES part of Poly1305-AES can be reused; the non-AES part of Poly1305-AES cannot be broken.
- **Extremely high speed.** Poly1305-AES can take advantage of the fast floating-point multipliers in popular CPUs. My Poly1305-AES software takes just 4212 Athlon cycles for a 1024-byte message. Special-purpose hardware can compute Poly1305-AES at even higher speed using integer multipliers. Much smaller processors can compute Poly1305-AES at reasonable speed using table lookups.
- **Low per-message overhead.** Poly1305-AES achieves good performance even for short messages. The same software takes just 577 Athlon cycles for a 16-byte message and 764 Athlon cycles for a 64-byte message. Most competing functions have much larger overhead for each message; they are optimized solely for large packets, even though a large fraction of Internet packets are below 64 bytes.
- **Key agility.** These cycle counts were achieved *with no precomputation.* The input is exactly (k, r, n, m) as stated above; the key (k, r) occupies exactly 32 bytes. Poly1305-AES authenticates messages very quickly even when it is cycling through 1000 simultaneous keys. This was my primary design goal for Poly1305-AES. Almost all competing functions use a large table for each key, typically several kilobytes per key, so they slow down dramatically (missing L1 cache and then missing L2 cache) as the number of keys grows; they are optimized solely for single-key applications.
- **Parallelizability and incrementality.** The circuit depth of Poly1305-AES is quite small, even for long messages. Consequently, Poly1305-AES can take advantage of additional hardware to reduce the latency for long messages. For essentially the same reason, Poly1305-AES can be recomputed at low cost for a small modification of a long message.
- **No intellectual-property claims.** I am not aware of any patents or patent applications relevant to Poly1305-AES.

Section 3 of this paper analyzes the security of Poly1305-AES. Section 4 discusses the software achieving the Athlon results stated above. Section 5 discusses the speed of Poly1305-AES in other contexts.

Genealogy

Gilbert, MacWilliams, and Sloane in [13] introduced the idea of provably secure authentication. The Gilbert-MacWilliams-Sloane system is fast, but it requires keys longer than L bytes to handle L-byte messages, and it requires a completely new key for each message.

Wegman and Carter in [30] pointed out that the key length could be merely $64 \lg L$ for the first message plus 16 bytes for each additional message. At about the same time, in a slightly different context, Karp and Rabin achieved a key length of 32 bytes for the first message; see [17] and [24]. The system in [17] is fast once keys are generated, but key generation is slow.

The idea of using a cipher such as AES to expand a short key into a long key is now considered obvious. Brassard in [10] published the idea in the Wegman-Carter context; I don't know whether the idea was considered obvious back then.

Polynomial-evaluation MACs—MACs that treat each message as a univariate polynomial over a finite field and then evaluate that polynomial at the key—were introduced in three papers independently: [12] by den Boer; [29, Section 3] by Taylor; [7, Section 4] by Bierbrauer, Johansson, Kabatianskii, and Smeets. Polynomial-evaluation MACs combine several attractive features: short keys, fast key generation, and fast message authentication. Several subsequent papers reported implementations of polynomial-evaluation MACs over binary fields: [26] by Shoup; [3] by Afanassiev, Gehrmann, and Smeets, reinventing Kaminski's division algorithm in [16]; [20] by Nevelsteen and Preneel.

Polynomial-evaluation MACs over prime fields can exploit the multipliers built into many current CPUs, achieving substantially better performance than polynomial-evaluation MACs over binary fields. This idea was first published in my paper [4] in April 1999, and explained in detail in [5]. Another MAC, avoiding binary fields for the same reason, was published independently by Black, Halevi, Krawczyk, Krovetz, and Rogaway in [9] in August 1999.

I used 32-bit polynomial coefficients modulo $2^{127} - 1$ ("hash127") in [4] and [5]. The short coefficients don't allow great performance (for short messages) without precomputation, so I casually precomputed a few kilobytes of data for each key; this is a disaster for applications handling many keys simultaneously, but I didn't think beyond a single key. Similarly, [9] ("UMAC") uses large keys.

Krovetz and Rogaway in [19] suggested 64-bit coefficients modulo $2^{64} - 59$, with an escape mechanism for coefficients between $2^{64} - 59$ and $2^{64} - 1$. They did not claim competitive performance: their software, run twice to achieve a reasonable 100-bit security level, was more than three times slower than hash127 (and more than six times slower for messages with all bits set). Krovetz and Rogaway did point out, however, that their software did not require large tables.

In `http://cr.yp.to/talks.html#2002.06.15`, posted July 2002, I pointed out that 128-bit coefficients over the slightly *larger* prime field $\mathbf{Z}/(2^{130} - 5)$ allow excellent performance without precomputation. This paper explains Poly1305-AES in much more detail.

Kohno, Viega, and Whiting subsequently suggested 96-bit coefficients modulo $2^{127} - 1$ ("CWC HASH"). They published some non-competitive timings for CWC HASH and then gave up on the idea. A careful implementation of CWC HASH without precomputation would be quite fast, although still not as fast as Poly1305-AES.

2 Specification

This section defines the Poly1305-AES function. The Poly1305-AES formula is a straightforward polynomial evaluation modulo $2^{130} - 5$; most of the detail is in key format and message padding.

Messages

Poly1305-AES authenticates messages. A **message** is any sequence of bytes $m[0], m[1], \ldots, m[\ell - 1] \in \{0, 1, \ldots, 255\}$. The message length ℓ can be any non-negative integer, and can vary from one message to another.

Keys

Poly1305-AES authenticates messages using a 32-byte secret key shared by the message sender and the message receiver. The key has two parts: first, a 16-byte AES key k; second, a 16-byte string $r[0], r[1], \ldots, r[15]$ representing a 128-bit integer r in unsigned little-endian form.

Certain bits of r are required to be 0: bytes $r[3], r[7], r[11], r[15]$ are required to have their top four bits clear, and bytes $r[4], r[8], r[12]$ are required to have their bottom two bits clear. In other words, r is required to have the form $r_0 + r_1 + r_2 + r_3$ where $r_0 \in \{0, 1, 2, 3, \ldots, 2^{28} - 1\}$, $r_1/2^{32} \in \{0, 4, 8, 12, \ldots, 2^{28} - 4\}$, $r_2/2^{64} \in \{0, 4, 8, 12, \ldots, 2^{28} - 4\}$, and $r_3/2^{96} \in \{0, 4, 8, 12, \ldots, 2^{28} - 4\}$. Thus there are 2^{106} possibilities for r.

Nonces

Poly1305-AES requires each message to be accompanied by a 16-byte **nonce**, i.e., a unique message number. Poly1305-AES feeds each nonce n through AES_k to obtain the 16-byte string $\mathrm{AES}_k(n)$.

There is nothing special about AES here. One can replace AES with an arbitrary keyed function from an arbitrary set of nonces to 16-byte strings. This paper focuses on AES for concreteness.

Conversion and padding

Let $m[0], m[1], \ldots, m[\ell - 1]$ be a message. Write $q = \lceil \ell/16 \rceil$. Define integers $c_1, c_2, \ldots, c_q \in \{1, 2, 3, \ldots, 2^{129}\}$ as follows: if $1 \le i \le \lfloor \ell/16 \rfloor$ then

$$c_i = m[16i - 16] + 2^8 m[16i - 15] + 2^{16} m[16i - 14] + \cdots + 2^{120} m[16i - 1] + 2^{128};$$

if ℓ is not a multiple of 16 then

$$c_q = m[16q - 16] + 2^8 m[16q - 15] + \cdots + 2^{8(\ell \bmod 16) - 8} m[\ell - 1] + 2^{8(\ell \bmod 16)}.$$

In other words: Pad each 16-byte chunk of a message to 17 bytes by appending a 1. If the message has a final chunk between 1 and 15 bytes, append 1 to the chunk, and then zero-pad the chunk to 17 bytes. Either way, treat the resulting 17-byte chunk as an unsigned little-endian integer.

Authenticators

$\text{Poly1305}_r(m, \text{AES}_k(n))$, the Poly1305-AES authenticator of a message m with nonce n under secret key (k, r), is defined as the 16-byte unsigned little-endian representation of

$$(((c_1 r^q + c_2 r^{q-1} + \cdots + c_q r^1) \bmod 2^{130} - 5) + \text{AES}_k(n)) \bmod 2^{128}.$$

Here the 16-byte string $\text{AES}_k(n)$ is treated as an unsigned little-endian integer, and $c_1, c_2, \ldots, c_q$ are the integers defined above. See Appendix A for examples.

Sample code

The following C++ code reads k from `k[0]`, `k[1]`, ..., `k[15]`, reads r from `r[0]`, `r[1]`, ..., `r[15]`, reads n from `n[0]`, `n[1]`, ..., `n[15]`, reads m from `m[0]`, `m[1]`, ..., `m[l-1]`, and places $\text{Poly1305}_r(m, \text{AES}_k(n))$ into `out[0]`, `out[1]`, ..., `out[15]`:

```
void poly1305aes_sample(unsigned char *out,
  const unsigned char *k,const unsigned char *r,
  const unsigned char *n,const unsigned char *m,unsigned int l)
{
  unsigned int j; integer rinteger = 0;
  for (j = 0;j < 16;++j)
    rinteger += ((integer) r[j]) << (8 * j);
  integer h = 0;
  integer p = (((integer) 1) << 130) - 5;
  while (l > 0) {
    integer c = 0;
    for (j = 0;(j < 16) && (j < l);++j)
      c += ((integer) m[j]) << (8 * j);
    c += ((integer) 1) << (8 * j);
    m += j; l -= j; h = ((h + c) * rinteger) % p;
  }
  unsigned char aeskn[16]; aes(aeskn,k,n);
  for (j = 0;j < 16;++j)
    h += ((integer) aeskn[j]) << (8 * j);
  for (j = 0;j < 16;++j) {
    integer c = h % 256; h >>= 8;
    out[j] = c.get_ui();
  }
}
```

Here `integer` can be defined as `mpz_class` from [14], or as any similar big-integer class.

This code is not meant as a high-speed implementation; it does not have even the simplest speedups; it should be expected to provide intolerable performance. It is simply a secondary statement of the definition of Poly1305-AES.

Design decisions

I considered various primes above 2^{128}. I chose $2^{130} - 5$ because its sparse form makes divisions particularly easy in both software and hardware. My encoding of messages as polynomials takes advantage of the gap between 2^{128} and $2^{130} - 5$.

There are several reasons that Poly1305-AES uses nonces. First, comparable protocols without nonces have security bounds that look like $C(C + D)L/2^{106}$ rather than $DL/2^{106}$—here C is the number of messages authenticated by the sender, D is the number of forgery attempts, and L is the maximum message length—and thus cannot be used with confidence for large C. Second, nonces allow the invocation of AES to be carried out in parallel with most of the other operations in Poly1305-AES, reducing latency in many contexts. Third, most protocols have nonces anyway, for a variety of reasons: nonces are required for secure encryption, for example, and nonces allow trivial rejection of replayed messages.

I constrained r to simplify and accelerate implementations of Poly1305-AES in various contexts. A wider range of r—e.g., all 128-bit integers—would allow a quantitatively better security bound, but the current bound $DL/2^{106}$ will be perfectly satisfactory for the foreseeable future, whereas slower authenticator computations would not be perfectly satisfactory.

I chose little-endian instead of big-endian to improve overall performance. Little-endian saves time on the most popular CPUs (the Pentium and Athlon) while making no difference on most other CPUs (the PowerPC, for example, and the UltraSPARC).

The definition of Poly1305-AES could easily be extended from byte strings to bit strings, but there is no apparent benefit of doing so.

3 Security

This section discusses the security of Poly1305-AES.

Responsibilities of the user

Any protocol that uses Poly1305-AES must ensure unpredictability of the secret key (k, r). This section assumes that secret keys are chosen from the uniform distribution: i.e., probability 2^{-234} for each of the 2^{234} possible pairs (k, r).

Any protocol that uses Poly1305-AES must ensure that the secret key is, in fact, kept secret. This section assumes that all operations are independent of (k, r), except for the computation of authenticators by the sender and receiver. (There are safe ways to reuse k for encryption, but those ways are not analyzed in this paper.)

The sender must *never* use the same nonce for two different messages. The simplest way to achieve this is for the sender to use an increasing sequence of nonces in, e.g., reverse-lexicographic order of 16-byte strings. (Problem: If a key is stored on disk, while increasing nonce values are stored in memory, what

happens when the power goes out? Solution: Store a safe nonce value—a new nonce larger than any nonce used—on disk alongside the key.) Any protocol that uses Poly1305-AES must specify a mechanism of nonce generation and maintenance that prevents duplicates.

Security guarantee

Poly1305-AES guarantees that the only way for the attacker to find an (n, m, a) such that $a = \text{Poly1305}_r(m, \text{AES}_k(n))$, other than the authenticated messages (n, m, a) sent by the sender, is to break AES. If the attacker cannot break AES, and the receiver discards all (n, m, a) such that $a \neq \text{Poly1305}_r(m, \text{AES}_k(n))$, then the receiver will see only messages authenticated by the sender.

This guarantee is not limited to "meaningful" messages m. It is true even if the attacker can see all the authenticated messages sent by the sender. It is true even if the attacker can see whether the receiver accepts a forgery. It is true even if the attacker can influence the sender's choice of messages and unique nonces. (But it is not true if the nonce-uniqueness rule is violated.)

Here is a quantitative form of the guarantee. Assume that the attacker sees at most C authenticated messages and attempts at most D forgeries. Assume that the attacker has probability at most δ of distinguishing AES_k from a uniform random permutation after $C + D$ queries. Assume that all messages have length at most L. Then, with probability at least

$$1 - \delta - \frac{(1 - C/2^{128})^{-(C+1)/2} 8D\lceil L/16 \rceil}{2^{106}},$$

all of the attacker's forgeries are discarded. In particular, if $C \leq 2^{64}$, then the attacker's chance of success is at most $\delta + 1.649 \cdot 8D\lceil L/16 \rceil / 2^{106} < \delta + 14D\lceil L/16 \rceil / 2^{106}$.

The most important design goal of AES was for δ to be small. There is, however, no hope of *proving* that δ is small. Perhaps AES will be broken someday. If that happens, users should switch to Poly1305-AnotherFunction. Poly1305-AnotherFunction provides the same security guarantee relative to the security of AnotherFunction.

Proof of the security guarantee

For each message m, write $\underline{m}$ for the polynomial $c_1 x^q + c_2 x^{q-1} + \cdots + c_q x^1$, where $q, c_1, c_2, \ldots, c_q$ are defined as in Section 2. Define $H_r(m)$ as the 16-byte unsigned little-endian representation of $(\underline{m}(r) \bmod 2^{130} - 5) \bmod 2^{128}$; note that H_r and k are independent. Define a group operation $+$ on 16-byte strings as addition modulo 2^{128}, where each 16-byte string is viewed as the unsigned little-endian representation of an integer in $\{0, 1, 2, \ldots, 2^{128} - 1\}$. Then the authenticator $\text{Poly1305}_r(m, \text{AES}_k(n))$ is equal to $H_r(m) + \text{AES}_k(n)$.

The crucial property of H_r is that it has small differential probabilities: if g is a 16-byte string, and m, m' are distinct messages of length at most L, then $H_r(m) = H_r(m') + g$ with probability at most $8\lceil L/16 \rceil / 2^{106}$. See below.

Theorem 5.4 of [6] now guarantees that $H_r(m) + \mathrm{AES}_k(n)$ is secure if AES is secure: specifically, that the attacker's success chance against $H_r(m) + \mathrm{AES}_k(n)$ is at most $\delta + D(1 - C/2^{128})^{-(C+1)/2} 8\lceil L/16 \rceil / 2^{106}$.

The rest of this section is devoted to proving that H_r has small differential probabilities.

Theorem 3.1. $2^{130} - 5$ *is prime.*

Proof. Define $p_1 = (2^{130} - 6)/1517314646$ and $p_2 = (p_1 - 1)/222890620702$. Observe that 37003 and 221101 are prime divisors of $p_2 - 1$; $(37003 \cdot 221101)^2 > p_2$; $2^{p_2 - 1} - 1$ is divisible by p_2; $2^{(p_2-1)/37003} - 1$ and $2^{(p_2-1)/221101} - 1$ are coprime to p_2; $p_2^2 > p_1$; $2^{p_1 - 1} - 1$ is divisible by p_1; $2^{(p_1-1)/p_2} - 1$ is coprime to p_1; $p_1^2 > 2^{130} - 5$; $2^{2^{130}-6} - 1$ is divisible by $2^{130} - 5$; and $2^{(2^{130}-6)/p_1} - 1$ is coprime to $2^{130} - 5$. Hence p_2, p_1, and $2^{130} - 5$ are prime by Pocklington's theorem. □

Theorem 3.2. *Let m and m' be messages. Let u be an integer. If the polynomial $\underline{m'} - \underline{m} - u$ is zero modulo $2^{130} - 5$ then $m = m'$.*

Proof. Define $c_1, c_2, \ldots, c_q$ as above, and define $c'_1, c'_2, \ldots, c'_{q'}$ for m' similarly.

If $q > q'$ then the coefficient of x^q in $\underline{m'} - \underline{m}$ is $0 - c_1$. By construction c_1 is in $\{1, 2, 3, \ldots, 2^{129}\}$, so it is nonzero modulo $2^{130} - 5$; contradiction. Thus $q \leq q'$. Similarly $q \geq q'$. Hence $q = q'$.

If $i \in \{1, 2, \ldots, q\}$ then $c_i - c'_i$ is the coefficient of x^{q+1-i} in $\underline{m'} - \underline{m} - u$, which by hypothesis is divisible by $2^{130} - 5$. But $c_i - c'_i$ is between -2^{129} and 2^{129} by construction. Hence $c_i = c'_i$. In particular, $c_q = c'_q$.

Define ℓ as the number of bytes in m. Recall that $q = \lceil \ell/16 \rceil$; thus ℓ is between $16q - 15$ and $16q$. The exact value of ℓ is determined by q and c_q: it is $16q$ if $2^{128} \leq c_q$, $16q - 1$ if $2^{120} \leq c_q < 2^{121}$, $16q - 2$ if $2^{112} \leq c_q < 2^{113}$, ..., $16q - 15$ if $2^8 \leq c_q < 2^9$. Hence m' also has ℓ bytes.

Now consider any $j \in \{0, 1, \ldots, \ell - 1\}$. Write $i = \lfloor j/16 \rfloor + 1$; then $16i - 16 \leq j \leq 16i - 1$, and $1 \leq i \leq \lceil \ell/16 \rceil = q$, so $m[j] = \left\lfloor c_i / 2^{8(j-16i+16)} \right\rfloor \bmod 256 = \left\lfloor c'_i / 2^{8(j-16i+16)} \right\rfloor \bmod 256 = m'[j]$. Hence $m = m'$. □

Theorem 3.3. *Let m, m' be distinct messages, each having at most L bytes. Let g be a 16-byte string. Let R be a subset of $\{0, 1, \ldots, 2^{130} - 6\}$. Then there are at most $8\lceil L/16 \rceil$ integers $r \in R$ such that $H_r(m) = H_r(m') + g$.*

Consequently, if $\#R = 2^{106}$, and if r is a uniform random element of R, then $H_r(m) = H_r(m') + g$ with probability at most $8\lceil L/16 \rceil / 2^{106}$.

Proof. Define U as the set of integers in $[-2^{130} + 6, 2^{130} - 6]$ congruent to g modulo 2^{128}. Note that $\#U \leq 8$.

If $H_r(m) = H_r(m') + g$ then $(\underline{m'}(r) \bmod 2^{130} - 5) - (\underline{m}(r) \bmod 2^{130} - 5) \equiv g \pmod{2^{128}}$ so $(\underline{m'}(r) \bmod 2^{130} - 5) - (\underline{m}(r) \bmod 2^{130} - 5) = u$ for some $u \in U$. Hence r is a root of the polynomial $\underline{m'} - \underline{m} - u$ modulo the prime $2^{130} - 5$. This polynomial is nonzero by Theorem 3.2, and has degree at most $\lceil L/16 \rceil$, so it has at most $\lceil L/16 \rceil$ roots modulo $2^{130} - 5$. Sum over all $u \in U$: there are most $8\lceil L/16 \rceil$ possibilities for r. □

4 A floating-point implementation

This section explains how to compute $\text{Poly1305}_r(m, \text{AES}_k(n))$, given (k, r, n, m), at very high speeds on the common Athlon CPU.

My software takes 577 cycles for a 16-byte message, 764 cycles for a 64-byte message, 1444 cycles for a 256-byte message, and 4212 cycles for a 1024-byte message—if the software is in L1 cache, (k, r, n, m) are in L1 cache, and (k, r, n, m) have memory locations divisible by 4. A more comprehensive speed table appears in Appendix B.

Outline

The overall strategy to compute $\text{Poly1305}_r(m, \text{AES}_k(n))$ is as follows. Start by setting an accumulator h to 0. For each chunk c of the message m, first set $h \leftarrow h + c$, and then set $h \leftarrow rh$. Periodically reduce h modulo $2^{130} - 5$, not necessarily to the smallest remainder but to something small enough to continue the computation. After all input chunks c are processed, fully reduce h modulo $2^{130} - 5$, and add $\text{AES}_k(n)$.

Large-integer arithmetic in floating-point registers

Represent each of h, c, r as a sum of floating-point numbers, as in [5]. Specifically:

- As in Section 2, write r as $r_0 + r_1 + r_2 + r_3$ where $r_0 \in \{0, 1, 2, \ldots, 2^{28} - 1\}$, $r_1/2^{32} \in \{0, 4, 8, \ldots, 2^{28} - 4\}$, $r_2/2^{64} \in \{0, 4, 8, \ldots, 2^{28} - 4\}$, and $r_3/2^{96} \in \{0, 4, 8, \ldots, 2^{28} - 4\}$. Store each of r_0, r_1, r_2, r_3, $5 \cdot 2^{-130} r_1$, $5 \cdot 2^{-130} r_2$, $5 \cdot 2^{-130} r_3$ in memory in 8-byte floating-point format.
- Write each message chunk c as $d_0 + d_1 + d_2 + d_3$ where $d_0, d_1/2^{32}, d_2/2^{64} \in \{0, 1, 2, 3, \ldots, 2^{32} - 1\}$ and $d_3/2^{96} \in \{0, 1, 2, 3, \ldots, 2^{34} - 1\}$.
- Write h as $h_0 + h_1 + h_2 + h_3$ where h_i is a multiple of 2^{32i} in the range specified below. Store each h_i in one of the Athlon's floating-point registers.

Warning: The FreeBSD operating system starts each program by instructing the CPU to round all floating-point mantissas to 53 bits, rather than using the CPU's natural 64-bit precision. Make sure to disable this instruction. Under `gcc`, for example, the code `asm volatile("fldcw %0"::"m"(0x137f))` specifies full 64-bit mantissas.

To set $h \leftarrow h + c$, set $h_0 \leftarrow h_0 + d_0$, $h_1 \leftarrow h_1 + d_1$, $h_2 \leftarrow h_2 + d_2$, $h_3 \leftarrow h_3 + d_3$. Before these additions, $h_0, h_1/2^{32}, h_2/2^{64}, h_3/2^{96}$ are required to be integers in $[-(63/128) \cdot 2^{64}, (63/128) \cdot 2^{64}]$. After these additions, $h_0, h_1/2^{32}, h_2/2^{64}, h_3/2^{96}$ are integers in $[-(127/256) \cdot 2^{64}, (127/256) \cdot 2^{64}]$.

Before multiplying h by r, reduce the range of each h_i by performing four parallel carries as follows. Define $\alpha_0 = 2^{95} + 2^{94}$, $\alpha_1 = 2^{127} + 2^{126}$, $\alpha_2 = 2^{159} + 2^{158}$, and $\alpha_3 = 2^{193} + 2^{192}$. Compute $y_i = \text{fp}_{64}(h_i + \alpha_i) - \alpha_i$ and $x_i = h_i - y_i$, where $\text{fp}_{64}(h_i + \alpha_i)$ means the 64-bit floating-point number closest to $h_i + \alpha_i$, with ties broken in the usual way. Then $y_0/2^{32}, y_1/2^{64}, y_2/2^{96}, y_3/2^{130}$ are integers. Set

$h_0 \leftarrow x_0 + 5 \cdot 2^{-130} y_3$, $h_1 \leftarrow x_1 + y_0$, $h_2 \leftarrow x_2 + y_1$, and $h_3 \leftarrow x_3 + y_2$. This substitution changes h by $(2^{130}-5)2^{-130}y_3$, so it does not change h mod $2^{130}-5$. There are 17 floating-point operations here: 8 additions, 8 subtractions, and 1 multiplication by the constant $5 \cdot 2^{-130}$.

Ranges: x_0, $x_1/2^{32}$, and $x_2/2^{64}$ are in $[-(1/2) \cdot 2^{32}, (1/2) \cdot 2^{32}]$; $x_3/2^{96}$ is in $[-2 \cdot 2^{32}, 2 \cdot 2^{32}]$; $y_0/2^{32}$, $y_1/2^{64}$, $y_2/2^{96}$, and $y_3/2^{128}$ are in $[-(127/256) \cdot 2^{32}, (127/256) \cdot 2^{32}]$; h_0 is in $[-(1147/1024) \cdot 2^{32}, (1147/1024) \cdot 2^{32}]$; $h_1/2^{32}$ is in $[-(255/256) \cdot 2^{32}, (255/256) \cdot 2^{32}]$; $h_2/2^{64}$ is in $[-(255/256) \cdot 2^{32}, (255/256) \cdot 2^{32}]$; $h_3/2^{96}$ is in $[-(639/256) \cdot 2^{32}, (639/256) \cdot 2^{32}]$.

To multiply h by r modulo $2^{130} - 5$, replace (h_0, h_1, h_2, h_3) with

$$\begin{aligned}
&(r_0h_0 + 5 \cdot 2^{-130} r_1h_3 + 5 \cdot 2^{-130} r_2h_2 + 5 \cdot 2^{-130} r_3h_1,\\
&\ r_0h_1 + \quad r_1h_0 + 5 \cdot 2^{-130} r_2h_3 + 5 \cdot 2^{-130} r_3h_2,\\
&\ r_0h_2 + \quad r_1h_1 + \quad r_2h_0 + 5 \cdot 2^{-130} r_3h_3,\\
&\ r_0h_3 + \quad r_1h_2 + \quad r_2h_1 + \quad r_3h_0).
\end{aligned}$$

Recall that $2^{-34}r_1$, $2^{-66}r_2$, and $2^{-98}r_3$ are integers, so $2^{-130}r_1h_3$, $2^{-130}r_2h_2$, and $2^{-130}r_3h_1$ are integers; similarly, $2^{-130}r_2h_3$ and $2^{-130}r_3h_2$ is a multiple of 2^{32}, and $2^{-130}r_3h_3$ is a multiple of 2^{64}. There are 28 floating-point operations here: 16 multiplications and 12 additions.

Ranges: $h_0, h_1/2^{32}, h_2/2^{64}, h_3/2^{96}$ are now integers of absolute value at most $2^{28}(1147/1024 + 2 \cdot (5/4)255/256 + (5/4)639/256)2^{32} < (63/128)2^{64}$, ready for the next iteration of the inner loop.

Note that the carries can be omitted on the first loop: d_0 is an integer in $[0, 2^{32}]$; $d_1/2^{32}$ is an integer in $[0, 2^{32}]$; $d_2/2^{64}$ is an integer in $[0, 2^{32}]$; $d_3/2^{96}$ is an integer in $[0, 3 \cdot 2^{32}]$; and $2^{28}(1 + (5/4) + (5/4) + (5/4)3)2^{32} < (63/128)2^{64}$.

Output conversion

After the last message chunk is processed, carry one last time, to put h_0, h_1, h_2, h_3 into the small ranges listed above.

Add $2^{130} - 2^{97}$ to h_3; add $2^{97} - 2^{65}$ to h_2; add $2^{65} - 2^{33}$ to h_1; and add $2^{33} - 5$ to h_0. This makes each h_i positive, and puts $h = h_0 + h_1 + h_2 + h_3$ into the range $\{0, 1, \ldots, 2(2^{130} - 5) - 1\}$.

Perform a few integer add-with-carry operations to convert the accumulator into a series of 32-bit words in the usual form. Subtract $2^{130} - 5$, and keep the result if it is nonnegative, being careful to use constant-time operations so that no information is leaked through timing.

Finally, add $\mathrm{AES}_k(n)$. My not-seriously-optimized Athlon implementation of AES takes approximately 102 cycles to expand k into a 176-byte table and 244 cycles to finish computing $\mathrm{AES}_k(n)$, if 8192 bytes of AES constants are in level-1 cache. Of course, a faster AES implementation will improve Poly1305-AES performance; the 102-cycle expansion can be eliminated if all active 176-byte expanded keys fit into level-1 cache; and a sufficiently advanced cache-loading mechanism would allow expanded keys to be fetched in the background during the computation of $\underline{m}(r)$.

Instruction selection and scheduling

Consider an integer (such as d_0) between 0 and $2^{32} - 1$, stored in the usual way as four bytes. How does one load the integer into a floating-point register, when the Athlon does not have a load-four-byte-unsigned-integer instruction? Here are three possibilities:

- Concatenate the four bytes with $(0, 0, 0, 0)$, and use the Athlon's load-eight-byte-signed-integer instruction. Unfortunately, the four-byte store forces the eight-byte load to wait for dozens of cycles.
- Concatenate the bytes with $(0, 0, 56, 67)$, producing an eight-byte floating-point number. Load that number, and subtract $2^{52}+2^{51}$ to obtain the desired integer. This well-known trick has the virtue of also allowing the integer to be scaled by (e.g.) 2^{32}: replace 67 with 69 and $2^{52} + 2^{51}$ with $2^{84} + 2^{83}$. Unfortunately, as above, the four-byte store forces the eight-byte load to wait for dozens of cycles.
- Subtract 2^{31} from the integer, use the Athlon's load-four-byte-signed-integer instruction, and add 2^{31} to the result. This is reasonably fast.

Top performance requires making the right choice.

(A variant of Poly1305-AES using signed 32-bit integers would save time on the Athlon. On the other hand, it would lose time on typical 64-bit CPUs.)

This is merely one example of several low-level issues that can drastically affect speed: instruction selection, instruction scheduling, register assignment, instruction fetching, etc. A "fast" implementation of Poly1305-AES, with just a few typical low-level mistakes, will use twice as many cycles per byte as the software described here.

Other modern CPUs

The software described above also works on the Pentium I, Pentium MMX, Pentium Pro, Pentium II, Pentium III, Pentium IV, and Pentium M. All of these chips have reasonably fast floating-point units.

The UltraSPARC, PowerPC, et al. support fast arithmetic on floating-point numbers with 53-bit, rather than 64-bit, mantissas. The simplest way to achieve good performance on these chips is to break a 32-bit number into two 16-bit pieces before multiplying it by another 32-bit number.

As above, careful attention to low-level CPU details is necessary for top performance. I plan to publish Poly1305-AES software for all of these CPUs.

5 Other implementation strategies

Some people, upon hearing that there is a tricky way to use the Athlon's floating-point unit to compute a function quickly, leap to the unjustified conclusion that the same function cannot be computed quickly except on an Athlon. Consider, for example, the incorrect statement "hash-127 needs good hardware support for a fast implementation" in [15, footnote 3].

This section outlines three non-floating-point methods to compute Poly1305-AES, and indicates contexts where the methods are useful.

Integer registers

The 130-bit accumulator in Poly1305-AES can be spread among several integer registers rather than several floating-point registers.

This is good for low-end CPUs that do not support floating-point operations but that still have reasonably fast integer multipliers. It is also good for some high-end CPUs, such as the Athlon 64, that offer faster multiplication through integer registers than through floating-point registers.

Tables

One can make a table of the integers $r, 2r, 4r, 8r, \ldots, 2^{129}r$ modulo $2^{130} - 5$, and then multiply any 130-bit integer by r by adding, on average, about 65 elements of the table.

One can reduce the amount of work by using both additions and subtractions, by increasing the table size, and by choosing table entries more carefully. For example, one can include $3r, 24r, 192r, \ldots$ in the table, and then multiply any 130-bit integer by r by adding and subtracting, on average, about 38 elements of the table. This is a special case of an algorithm often credited to Brickell, Gordon, McCurley, Wilson, Lim, and Lee, but actually introduced much earlier by Pippenger in [21].

One can also balance the table size against the effort in reduction modulo $2^{130} - 5$. Consider, for example, the table $r, 2r, 3r, 4r, \ldots, 255r$.

Table lookups are often the best approach for tiny CPUs that do not have any fast multiplication operations. Of course, their key agility is poor, and they are susceptible to timing attacks if they are not implemented very carefully.

Special-purpose circuits

A 1600MHz AMD Duron, costing under $50, can feed 3 gigabits per second of 1500-byte messages through Poly1305-AES with the software discussed in Section 4. Hardware implementations of Poly1305-AES can strip away a great deal of unnecessary cost: the multiplier is only part of the cost of the Duron; furthermore, some of the multiplications are by sparse constants; furthermore, only about 20% of the multiplier area is doing any useful work, since each input is much smaller than 64 bits; furthermore, almost all carries can be deferred until the end of the Poly1305-AES computation, rather than being performed after each multiplication; furthermore, hardware implementations need not, and should not, imitate traditional software structures—one can directly build a fast multiplier modulo $2^{130}-5$, taking advantage of more sophisticated multiplication algorithms than those used in the Duron. Evidently Poly1305-AES can handle next-generation Ethernet speeds at reasonable cost.

References

1. —, *17th annual symposium on foundations of computer science*, IEEE Computer Society, Long Beach, California, 1976. MR 56:1766.
2. —, *20th annual symposium on foundations of computer science*, IEEE Computer Society, New York, 1979. MR 82a:68004.
3. Valentine Afanassiev, Christian Gehrmann, Ben Smeets, *Fast message authentication using efficient polynomial evaluation*, in [8] (1997), 190–204. URL: `http://cr.yp.to/bib/entries.html#1997/afanassiev`.
4. Daniel J. Bernstein, *Guaranteed message authentication faster than MD5 (abstract)* (1999). URL: `http://cr.yp.to/papers.html#hash127-abs`.
5. Daniel J. Bernstein, *Floating-point arithmetic and message authentication*, to be incorporated into author's *High-speed cryptography* book. URL: `http://cr.yp.to/papers.html#hash127`. ID `dabadd3095644704c5cbe9690ea3738e`.
6. Daniel J. Bernstein, *Stronger security bounds for Wegman-Carter-Shoup authenticators*. URL: `http://cr.yp.to/papers.html#securitywcs`. ID `2d603727f69542f30f7da2832240c1ad`.
7. Jürgen Bierbrauer, Thomas Johansson, Gregory Kabatianskii, Ben Smeets, *On families of hash functions via geometric codes and concatenation*, in [28] (1994), 331–342. URL: `http://cr.yp.to/bib/entries.html#1994/bierbrauer`.
8. Eli Biham (editor), *Fast Software Encryption '97*, Lecture Notes in Computer Science, 1267, Springer-Verlag, Berlin, 1997. ISBN 3–540–63247–6.
9. John Black, Shai Halevi, Hugo Krawczyk, Ted Krovetz, Phillip Rogaway, *UMAC: fast and secure message authentication*, in [32] (1999), 216–233. URL: `http://www.cs.ucdavis.edu/~rogaway/umac/`.
10. Gilles Brassard, *On computationally secure authentication tags requiring short secret shared keys*, in [11] (1983), 79–86. URL: `http://cr.yp.to/bib/entries.html#1983/brassard`.
11. David Chaum, Ronald L. Rivest, Alan T. Sherman (editors), *Advances in cryptology: proceedings of Crypto 82*, Plenum Press, New York, 1983. ISBN 0–306–41366–3. MR 84j:94004.
12. Bert den Boer, *A simple and key-economical unconditional authentication scheme*, Journal of Computer Security **2** (1993), 65–71. ISSN 0926–227X. URL: `http://cr.yp.to/bib/entries.html#1993/denboer`.
13. Edgar N. Gilbert, F. Jessie MacWilliams, Neil J. A. Sloane, *Codes which detect deception*, Bell System Technical Journal **53** (1974), 405–424. ISSN 0005–8580. MR 55:5306. URL: `http://cr.yp.to/bib/entries.html#1974/gilbert`.
14. Torbjorn Granlund (editor), *GMP 4.1.2: GNU multiple precision arithmetic library* (2004). URL: `http://www.swox.com/gmp/`.
15. Shai Halevi, Phil Rogaway, *A tweakable enciphering mode* (2003). URL: `http://www.research.ibm.com/people/s/shaih/pubs/hr03.html`.
16. Michael Kaminski, *A linear time algorithm for residue computation and a fast algorithm for division with a sparse divisor*, Journal of the ACM **34** (1987), 968–984. ISSN 0004–5411. MR 89f:68033.
17. Richard M. Karp, Michael O. Rabin, *Efficient randomized pattern-matching algorithms*, IBM Journal of Research and Development **31** (1987), 249–260. ISSN 0018–8646. URL: `http://cr.yp.to/bib/entries.html#1987/karp`.
18. Neal Koblitz (editor), *Advances in cryptology—CRYPTO '96*, Lecture Notes in Computer Science, 1109, Springer-Verlag, Berlin, 1996.

19. Ted Krovetz, Phillip Rogaway, *Fast universal hashing with small keys and no preprocessing: the PolyR construction* (2000). URL: `http://www.cs.ucdavis.edu/~rogaway/papers/poly.htm`.
20. Wim Nevelsteen, Bart Preneel, *Software performance of universal hash functions*, in [27] (1999), 24–41.
21. Nicholas Pippenger, *On the evaluation of powers and related problems (preliminary version)*, in [1] (1976), 258–263; newer version split into [22] and [23]. MR 58:3682. URL: `http://cr.yp.to/bib/entries.html#1976/pippenger`.
22. Nicholas Pippenger, *The minimum number of edges in graphs with prescribed paths*, Mathematical Systems Theory **12** (1979), 325–346; see also older version [21]. ISSN 0025–5661. MR 81e:05079. URL: `http://cr.yp.to/bib/entries.html#1979/pippenger`.
23. Nicholas Pippenger, *On the evaluation of powers and monomials*, SIAM Journal on Computing **9** (1980), 230–250; see also older version [21]. ISSN 0097–5397. MR 82c:10064. URL: `http://cr.yp.to/bib/entries.html#1980/pippenger`.
24. Michael O. Rabin, *Fingerprinting by random polynomials*, Harvard Aiken Computational Laboratory TR-15-81 (1981). URL: `http://cr.yp.to/bib/entries.html#1981/rabin`.
25. Victor Shoup, *On fast and provably secure message authentication based on universal hashing*, in [18] (1996), 313–328; see also newer version [26].
26. Victor Shoup, *On fast and provably secure message authentication based on universal hashing* (1996); see also older version [25]. URL: `http://www.shoup.net/papers`.
27. Jacques Stern (editor), *Advances in cryptology: EUROCRYPT '99*, Lecture Notes in Computer Science, 1592, Springer-Verlag, Berlin, 1999. ISBN 3–540–65889–0. MR 2000i:94001.
28. Douglas R. Stinson (editor), *Advances in cryptology—CRYPTO '93: 13th annual international cryptology conference, Santa Barbara, California, USA, August 22–26, 1993, proceedings*, Lecture Notes in Computer Science, 773, Springer-Verlag, Berlin, 1994. ISBN 3–540–57766–1, 0–387–57766–1. MR 95b:94002.
29. Richard Taylor, *An integrity check value algorithm for stream ciphers*, in [28] (1994), 40–48. URL: `http://cr.yp.to/bib/entries.html#1994/taylor`.
30. Mark N. Wegman, J. Lawrence Carter, *New classes and applications of hash functions*, in [2] (1979), 175–182; see also newer version [31]. URL: `http://cr.yp.to/bib/entries.html#1979/wegman`.
31. Mark N. Wegman, J. Lawrence Carter, *New hash functions and their use in authentication and set equality*, Journal of Computer and System Sciences **22** (1981), 265–279; see also older version [30]. ISSN 0022–0000. MR 82i:68017. URL: `http://cr.yp.to/bib/entries.html#1981/wegman`.
32. Michael Wiener (editor), *Advances in cryptology—CRYPTO '99*, Lecture Notes in Computer Science, 1666, Springer-Verlag, Berlin, 1999. ISBN 3–5540–66347–9. MR 2000h:94003.

A Appendix: Examples

The following table, with all integers on the right displayed in hexadecimal, illustrates authenticator computations for strings of length 2, 0, 32, and 63. The notation $\underline{m}(r)$ means $c_1 r^q + c_2 r^{q-1} + \cdots + c_q r^1$.

m	f3 f6
c_1	00000000000000000000000000001f6f3
r	85 1f c4 0c 34 67 ac 0b e0 5c c2 04 04 f3 f7 00
$\underline{m}(r) \bmod 2^{130} - 5$	321e58e25a69d7f8f27060770b3f8bb9c
k	ec 07 4c 83 55 80 74 17 01 42 5b 62 32 35 ad d6
n	fb 44 73 50 c4 e8 68 c5 2a c3 27 5c f9 d4 32 7e
$\mathrm{AES}_k(n)$	58 0b 3b 0f 94 47 bb 1e 69 d0 95 b5 92 8b 6d bc
$\mathrm{Poly1305}_r(m, \mathrm{AES}_k(n))$	f4 c6 33 c3 04 4f c1 45 f8 4f 33 5c b8 19 53 de
m	
r	a0 f3 08 00 00 f4 64 00 d0 c7 e9 07 6c 83 44 03
$\underline{m}(r) \bmod 2^{130} - 5$	00000000000000000000000000000000
k	75 de aa 25 c0 9f 20 8e 1d c4 ce 6b 5c ad 3f bf
n	61 ee 09 21 8d 29 b0 aa ed 7e 15 4a 2c 55 09 cc
$\mathrm{AES}_k(n)$	dd 3f ab 22 51 f1 1a c7 59 f0 88 71 29 cc 2e e7
$\mathrm{Poly1305}_r(m, \mathrm{AES}_k(n))$	dd 3f ab 22 51 f1 1a c7 59 f0 88 71 29 cc 2e e7
m	66 3c ea 19 0f fb 83 d8 95 93 f3 f4 76 b6 bc 24
	d7 e6 79 10 7e a2 6a db 8c af 66 52 d0 65 61 36
c_1	124bcb676f4f39395d883fb0f19ea3c66
c_2	1366165d05266af8cdb6aa27e1079e6d7
r	48 44 3d 0b b0 d2 11 09 c8 9a 10 0b 5c e2 c2 08
$\underline{m}(r) \bmod 2^{130} - 5$	1cfb6f98add6a0ea7c631de020225cc8b
k	6a cb 5f 61 a7 17 6d d3 20 c5 c1 eb 2e dc dc 74
n	ae 21 2a 55 39 97 29 59 5d ea 45 8b c6 21 ff 0e
$\mathrm{AES}_k(n)$	83 14 9c 69 b5 61 dd 88 29 8a 17 98 b1 07 16 ef
$\mathrm{Poly1305}_r(m, \mathrm{AES}_k(n))$	0e e1 c1 6b b7 3f 0f 4f d1 98 81 75 3c 01 cd be
m	ab 08 12 72 4a 7f 1e 34 27 42 cb ed 37 4d 94 d1
	36 c6 b8 79 5d 45 b3 81 98 30 f2 c0 44 91 fa f0
	99 0c 62 e4 8b 80 18 b2 c3 e4 a0 fa 31 34 cb 67
	fa 83 e1 58 c9 94 d9 61 c4 cb 21 09 5c 1b f9
c_1	1d1944d37edcb4227341e7f4a721208ab
c_2	1f0fa9144c0f2309881b3455d79b8c636
c_3	167cb3431faa0e4c3b218808be4620c99
c_4	001f91b5c0921cbc461d994c958e183fa
r	12 97 6a 08 c4 42 6d 0c e8 a8 24 07 c4 f4 82 07
$\underline{m}(r) \bmod 2^{130} - 5$	0c3c4f37c464bbd44306c9f8502ea5bd1
k	e1 a5 66 8a 4d 5b 66 a5 f6 8c c5 42 4e d5 98 2d
n	9a e8 31 e7 43 97 8d 3a 23 52 7c 71 28 14 9e 3a
$\mathrm{AES}_k(n)$	80 f8 c2 0a a7 12 02 d1 e2 91 79 cb cb 55 5a 57
$\mathrm{Poly1305}_r(m, \mathrm{AES}_k(n))$	51 54 ad 0d 2c b2 6e 01 27 4f c5 11 48 49 1f 1b

B Appendix: Speed table

The table below shows cycle counts on a 900MHz Athlon for ten Poly1305-AES computations on a 0-byte message, then ten Poly1305-AES computations on a 1-byte message, etc., using the software described in Section 4.

Each number includes about 35 cycles of timing overhead and C function-call overhead. I obtained each set of ten numbers as consecutive differences of eleven results from the Athlon RDTSC instruction surrounding ten calls to the Poly1305-AES function.

There are several reasons for variability within lines:

- My Poly1305-AES code takes 4375 bytes (of which 3046 bytes are for AES). For message lengths 0, 1, 16, and 17, the first computation needed extra time to load the relevant code into L1 cache.
- The next few length-0 computations needed extra time to load unused parts of the 8192-byte AES table into L1 cache.
- The remaining length-0 computations needed extra time to load new keys into L1 cache. I used ten separate keys (k, r).
- Each new message length meant a new sequence of conditional jumps; the Athlon's branch-prediction unit took a few iterations to adapt to the new sequence.
- AES performance depended slightly on the key.

I'm writing a separate paper on the problem of eliminating load-timing leaks in fast AES software.

0	17329	3949	1553	736	927	612	735	768	416	422
1	3668	833	585	574	603	567	577	568	570	585
2	568	572	574	575	570	563	565	569	571	574
3	569	573	575	576	571	564	566	570	572	575
4	573	577	579	580	575	568	570	574	576	579
5	574	578	580	581	576	569	571	575	577	580
6	575	579	581	582	577	570	572	576	578	581
7	577	581	583	584	579	572	574	578	580	583
8	578	582	584	585	580	573	575	579	581	584
9	579	583	585	586	581	574	576	580	582	585
10	581	585	587	588	583	576	578	582	584	587
11	582	586	588	589	584	577	579	583	585	588
12	591	593	595	596	591	584	586	590	592	595
13	591	595	597	598	593	586	588	592	594	597
14	592	596	598	599	594	587	589	593	595	598
15	593	597	599	600	595	588	590	594	596	599
16	577	552	566	555	550	543	545	549	563	566
17	1149	657	659	644	639	632	634	638	640	643
18	646	658	644	645	640	633	635	639	641	644
19	646	643	645	646	641	634	636	640	642	645
20	643	647	649	650	645	638	640	644	646	649

21	644	648	650	651	646	639	641	645	647	650
22	645	649	651	652	647	640	642	646	648	651
23	647	651	653	654	649	642	644	648	650	653
24	652	656	658	659	654	647	649	653	655	658
25	653	657	659	660	655	648	650	654	656	659
26	655	659	661	662	657	650	652	656	658	661
27	656	660	662	663	658	651	653	657	659	662
28	659	663	665	666	661	654	656	660	662	665
29	661	665	667	668	663	656	658	662	664	667
30	662	666	668	669	664	657	659	663	665	668
31	663	667	669	670	665	658	660	664	666	669
32	654	653	665	630	610	603	605	609	611	614
33	729	743	719	703	698	691	693	697	699	702
34	713	701	703	704	699	692	694	698	700	703
35	697	702	704	705	700	693	695	699	701	704
36	702	706	708	709	704	697	699	703	705	708
37	703	707	709	710	705	698	700	704	706	709
38	704	708	710	711	706	699	701	705	707	710
39	706	710	712	713	708	701	703	707	709	712
40	711	715	717	718	713	706	708	712	714	717
41	728	716	718	719	714	707	709	713	715	718
42	714	718	720	721	716	709	711	715	717	720
43	715	719	721	722	717	710	712	716	718	721
44	714	719	721	722	717	710	712	716	718	721
45	716	721	723	724	719	712	714	718	720	723
46	717	722	724	725	720	713	715	719	721	724
47	718	723	725	726	721	714	716	720	722	725
48	711	724	691	673	668	661	663	667	669	672
64	758	762	764	731	726	719	721	725	727	730
80	816	834	822	788	783	776	778	782	784	787
96	883	877	863	846	841	834	836	840	842	845
128	998	992	960	961	956	949	951	955	957	960
160	1122	1136	1091	1092	1087	1080	1082	1086	1088	1091
192	1210	1214	1206	1207	1202	1195	1197	1201	1203	1206
224	1325	1329	1321	1322	1317	1310	1312	1316	1318	1321
256	1440	1444	1436	1437	1432	1425	1427	1431	1433	1436
320	1670	1674	1666	1667	1662	1655	1657	1661	1663	1666
384	1900	1904	1896	1897	1892	1885	1887	1891	1893	1896
448	2138	2142	2134	2135	2130	2123	2125	2129	2131	2134
512	2366	2372	2364	2365	2360	2353	2355	2359	2361	2364
768	3288	3292	3284	3285	3280	3273	3275	3279	3281	3284
1024	4206	4212	4204	4205	4200	4193	4195	4199	4201	4204
1536	6048	6052	6044	6045	6040	6033	6035	6039	6041	6044
2048	7888	7892	7884	7885	7880	7873	7875	7879	7881	7884
3072	11579	11572	11564	11565	11560	11553	11555	11559	11561	11564
4096	15259	15252	15239	15245	15235	15228	15235	15239	15241	15244