

Cryptography in NaCl

Daniel J. Bernstein *

Department of Computer Science (MC 152)
The University of Illinois at Chicago
Chicago, IL 60607-7053
djb@cr.yp.to

1 Introduction

“NaCl” (pronounced “salt”) is the CACE Networking and Cryptography library, a new easy-to-use high-speed high-security public-domain software library for network communication, encryption, decryption, signatures, etc. Of course, other libraries already exist for these core operations; NaCl advances the state of the art by improving security, by improving usability, and by improving speed.

The most fundamental operation in a cryptographically protected network protocol is **public-key authenticated encryption**. The sender, Alice, has a packet of data to send to the receiver, Bob. Alice scrambles the packet using her own secret key and Bob’s public key. Bob unscrambles the packet using his secret key and Alice’s public key. An attacker monitoring the network is unable to understand the scrambled packet; an attacker modifying network packets is unable to change the packet produced by Bob’s unscrambling.

With typical cryptographic libraries, public-key authenticated encryption takes several steps. Here is a typical series of steps:

- Generate a random AES key.
- Use the AES key to encrypt the packet.
- Hash the encrypted packet using SHA-256.
- Read Alice’s RSA secret key from “wire format.”
- Use Alice’s RSA secret key to sign the hash.
- Read Bob’s RSA public key from wire format.
- Use Bob’s public key to encrypt the AES key, hash, and signature.
- Convert the encrypted key, hash, and signature to wire format.
- Concatenate with the encrypted packet.

NaCl provides a high-level function `crypto_box` that does everything in one step, converting a packet into a boxed packet that is protected against espionage and sabotage. Programmers can use lower-level functions but are encouraged to use `crypto_box`.

In particular, `crypto_box_curve25519xsalsa20poly1305` is a specific high-speed high-security combination of the Curve25519 elliptic-curve-Diffie–Hellman

* Permanent ID of this document: `1ae6a0ecef3073622426b3ee56260d34`. Date of this document: 2009.03.10.

function, the Salsa20 stream cipher, and the Poly1305 message-authentication code. This combination is designed for universal use and is shipped in NaCl as the default definition of `crypto_box`.

This document specifies exactly what this combination does: i.e., exactly how the boxed packet produced by `crypto_box_curve25519xsalsa20poly1305` relates to the inputs. The specification is expressed as a step-by-step procedure for Alice to encrypt and authenticate a packet; NaCl might compute the boxed packet in a different way but produces exactly the same results. Three of the steps are packet-independent precomputation:

- Section 2: Alice creates a 32-byte secret key a and a 32-byte public key A . These keys can be reused for other packets to Bob, for packets to other receivers, and for packets sent back from the receivers.
- Section 2, continued: Bob creates a 32-byte secret key b and a 32-byte public key B . These keys can be reused for other packets from Alice, for packets from other senders, and for packets sent back to the senders.
- Section 5: Alice, using Alice's secret key a and Bob's public key B , computes a 32-byte secret k . Bob can compute the same secret using Bob's secret key b and Alice's public key A .

The remaining three steps are specific to one packet:

- Section 7: Alice, using a 24-byte nonce (unique packet number) n that will never be reused for other packets to (or from) Bob, expands the shared secret k into a long stream of secret bytes. Bob, given the nonce, can compute the same stream.
- Section 9: Alice uses the long stream, except for the first 32 bytes, to encrypt the packet m .
- Section 9, continued: Alice uses the first 32 bytes of the long stream to compute an authenticator of the encrypted packet.

Each section includes security notes and pointers to the relevant literature.

This document also contains, in Sections 3, 4, 6, 8, and 10, a complete step-by-step example to illustrate the specification. The intermediate results are printed by various C NaCl programs shown here.

This document also contains several tests showing that C NaCl is producing the same results as independent programs in other languages. Some of the tests rely on scripts using the Sage computer-algebra system [20], and some of the tests rely on Python scripts contributed by Matthew Dempsky. This document can be used as a starting point for more comprehensive NaCl validation and verification.

In this document, a **byte** is an element of $\{0, 1, \dots, 255\}$. NaCl works with all keys, packets, etc. as strings of bytes. For example, the set of 32-byte strings is the set $\{0, 1, \dots, 255\}^{32}$.

2 Secret keys and public keys

Alice's secret key is a string $a \in \{0, 1, \dots, 255\}^{32}$. Alice's public key is a string $\text{Curve25519}(\text{ClampC}(a), \underline{9}) \in \{0, 1, \dots, 255\}^{32}$. Similarly, Bob's secret key is a

string $b \in \{0, 1, \dots, 255\}^{32}$, and Bob's public key is $\text{Curve25519}(\text{ClampC}(b), \underline{9}) \in \{0, 1, \dots, 255\}^{32}$.

This section defines the functions `ClampC` and `Curve25519` and the constant `9`. Many of the definitions here are copied from [8, Section 2]; in particular, `Curve25519` here is the same as the `Curve25519` function defined in [8].

Section 3 gives an example of a secret key and corresponding public key that Alice might choose. Section 4 gives an example of a secret key and corresponding public key that Bob might choose. These examples are reused in subsequent sections.

The base field and the elliptic curve. Define $p = 2^{255} - 19$. This integer is prime:

```
sage: p=2^255-19
sage: p.is_prime()
True
```

Define \mathbf{F}_p as the prime field $\mathbf{Z}/p = \mathbf{Z}/(2^{255} - 19)$. Note that 2 is not a square in \mathbf{F}_p :

```
sage: p=2^255-19
sage: k=GF(p)
sage: k(2).is_square()
False
```

Define \mathbf{F}_{p^2} as the field $(\mathbf{Z}/(2^{255} - 19))[\sqrt{2}]$. Define $a_2 = 486662$. Note that $a_2^2 - 4$ is not a square in \mathbf{F}_p :

```
sage: p=2^255-19
sage: k=GF(p)
sage: a2=486662
sage: (k(a2)^2-4).is_square()
False
```

Define E as the elliptic curve $y^2 = x^3 + a_2x^2 + x$ over \mathbf{F}_p , and define $E(\mathbf{F}_{p^2})$ as the group of points of E with coordinates in \mathbf{F}_{p^2} . Readers not familiar with elliptic curves can find a self-contained definition of $E(\mathbf{F}_{p^2})$ in [8, Appendix A].

Define $X_0 : E(\mathbf{F}_{p^2}) \rightarrow \mathbf{F}_{p^2}$ as follows: $X_0(\infty) = 0$; $X_0(x, y) = x$.

The Curve25519 function. Write $s \mapsto \underline{s}$ for the standard little-endian bijection from $\{0, 1, \dots, 2^{256} - 1\}$ to $\{0, 1, \dots, 255\}^{32}$. In other words, for each integer $s \in \{0, 1, \dots, 2^{256} - 1\}$, define

$$\underline{s} = (s \bmod 256, \lfloor s/256 \rfloor \bmod 256, \dots, \lfloor s/256^{31} \rfloor \bmod 256).$$

For example, the constant `9` is $(9, 0, 0, \dots, 0) \in \{0, 1, \dots, 255\}^{32}$.

The set of **Curve25519 secret keys** is, by definition, $\{0, 8, 16, 24, \dots, 248\} \times \{0, 1, \dots, 255\}^{30} \times \{64, 65, 66, \dots, 127\}$. If $n \in 2^{254} + 8\{0, 1, 2, 3, \dots, 2^{251} - 1\}$

then \underline{n} is a Curve25519 secret key; and every Curve25519 secret key can be written as \underline{n} for a unique $n \in 2^{254} + 8\{0, 1, 2, 3, \dots, 2^{251} - 1\}$.

Now the function

$$\text{Curve25519} : \{\text{Curve25519 secret keys}\} \times \{0, 1, \dots, 255\}^{32} \rightarrow \{0, 1, \dots, 255\}^{32}$$

is defined as follows. Fix an integer $n \in 2^{254} + 8\{0, 1, 2, 3, \dots, 2^{251} - 1\}$ and an integer $q \in \{0, 1, \dots, 2^{256} - 1\}$. By [8, Theorem 2.1] there is a unique integer $s \in \{0, 1, \dots, 2^{255} - 20\}$ such that $X_0(nQ) = s$ for all $Q \in E(\mathbf{F}_{p^2})$ such that $X_0(Q) = q \bmod 2^{255} - 19$. Finally, $\text{Curve25519}(\underline{n}, \underline{q})$ is defined as \underline{s} .

The ClampC function. The function

$$\text{ClampC} : \{0, 1, \dots, 255\}^{32} \rightarrow \{\text{Curve25519 secret keys}\}$$

maps $(a_0, a_1, \dots, a_{30}, a_{31})$ to $(a_0 - (a_0 \bmod 8), a_1, \dots, a_{30}, 64 + (a_{31} \bmod 64))$. In other words, ClampC clears bits (7, 0, ..., 0, 0, 128) and sets bit (0, 0, ..., 0, 0, 64).

Specialization of Curve25519 for secret keys. Note that $9^3 + a_2 \cdot 9^2 + 9 = 39420360$:

```
sage: a2=486662
sage: 9^3+a2*9^2+9
39420360
```

If $n \in 2^{254} + 8\{0, 1, 2, 3, \dots, 2^{251} - 1\}$ then $\text{Curve25519}(\underline{n}, \underline{9}) = \underline{s}$, where s is the unique integer in $\{0, 1, \dots, 2^{255} - 20\}$ such that $X_0(n(9, \pm\sqrt{39420360})) = s$. Consequently, if Alice's secret key a satisfies $\text{ClampC}(a) = \underline{n}$, then Alice's public key is \underline{s} .

The range of n implies that $n(9, \pm\sqrt{39420360}) \neq \infty$, so ∞ could be omitted from the definition of X_0 for purposes of computing secret keys. However, Alice also applies Curve25519 to network inputs, as discussed in subsequent sections, and there are several ways that attacker-chosen inputs can lead to the ∞ case.

ECDLP security notes. The following notes assume additional familiarity with elliptic curves.

Write $Q = (9, \sqrt{39420360})$. The choice of square root is not relevant here. This point Q is in the subgroup $E(\mathbf{F}_p)$ of $E(\mathbf{F}_{p^2})$:

```
sage: p=2^255-19
sage: k=GF(p)
sage: k(39420360).is_square()
True
```

Furthermore, Q has p_1 th multiple equal to ∞ in $E(\mathbf{F}_p)$, where p_1 is the prime number $2^{252} + 27742317777372353535851937790883648493$:

```
sage: p=2^255-19
sage: k=GF(p)
sage: p1=2^252+27742317777372353535851937790883648493
```

```

sage: p1.is_prime()
True
sage: E=EllipticCurve([k(0),486662,0,1,0])
sage: Q=[k(9),sqrt(k(39420360))]
sage: p1*E(Q)
(0 : 1 : 0)

```

Consequently all multiples of Q are in the subgroup of $E(\mathbf{F}_p)$ of order p_1 .

If Alice's secret key a is a uniform random 32-byte string then $\text{ClampC}(a)$ is a uniform random Curve25519 secret key; i.e., \underline{n} , where $n/8$ is a uniform random integer between 2^{251} and $2^{252} - 1$. Alice's public key is nQ compressed to the x -coordinate (as recommended in [17, page 425, fourth paragraph] in 1986). Note that n is not a multiple of p_1 ; this justifies the statement above that $nQ \neq \infty$.

The problem of finding Alice's secret key from Alice's public key is exactly the elliptic-curve 251-bit-discrete-logarithm problem for the subgroup of $E(\mathbf{F}_p)$ of order $p_1 \approx 2^{252}$. The curve E meets all of the standard security criteria, as discussed in detail in [8, Section 3]. The fastest known attacks use, on average, about 2^{125} additions in $E(\mathbf{F}_p)$, and have success chance degrading quadratically as the number of additions decreases.

It is standard in the literature to restrict attention to uniform random secret keys. What if the key distribution is not uniform? The answer depends on the distribution. For example, a key derived from an 8-byte string can be found by brute-force search in roughly 2^{64} operations; and a key derived in an *extremely weak* way from a 16-byte string, for example by concatenating a 16-byte public constant, can also be found in roughly 2^{64} operations. On the other hand, it is easy to prove that *slightly* non-uniform keys have essentially full security. Furthermore, I am not aware of any feasible attacks against 32-byte keys of the form $(s, \text{MD5}(s))$, where s is a uniform random 16-byte string; none of the weaknesses of MD5 seem relevant here. Constructions of this type allow secret-key compression and might merit further study if there are any applications where memory is filled with secret keys.

3 Example of the sender's keys

The following program uses C NaCl to compute the public key corresponding to a particular secret key:

```

#include <stdio.h>
#include "crypto_scalarmult_curve25519.h"

unsigned char alicesk[32] = {
    0x77,0x07,0x6d,0x0a,0x73,0x18,0xa5,0x7d
    ,0x3c,0x16,0xc1,0x72,0x51,0xb2,0x66,0x45
    ,0xdf,0x4c,0x2f,0x87,0xeb,0xc0,0x99,0x2a
    ,0xb1,0x77,0xfb,0xa5,0x1d,0xb9,0x2c,0x2a
} ;

```

```

unsigned char alicepk[32];

main()
{
    int i;
    crypto_scalarmult_curve25519_base(alicepk,alicesk);
    for (i = 0;i < 32;++i) {
        if (i > 0) printf(","); else printf(" ");
        printf("0x%02x",(unsigned int) alicepk[i]);
        if (i % 8 == 7) printf("\n");
    }
    return 0;
}

```

The secret key bytes `0xc7,0x6e,...` embedded into the program were copied from output of `od -t x1 /dev/urandom | head -2`. The output of the program is the corresponding public key:

```

0x85,0x20,0xf0,0x09,0x89,0x30,0xa7,0x54
,0x74,0x8b,0x7d,0xdc,0xb4,0x3e,0xf7,0x5a
,0x0d,0xbf,0x3a,0x0d,0x26,0x38,0x1a,0xf4
,0xeb,0xa4,0xa9,0x8e,0xaa,0x9b,0x4e,0x6a

```

The remaining sections of this document will reuse this example, assuming that Alice's keys are the particular secret key and public key shown here.

Testing: Sage vs. `scalarmult_curve25519_base`. A short Sage script clamps Alice's secret key shown above, converts the result to an integer n , computes $n(9, \sqrt{39420360})$ in $E(\mathbf{F}_p)$, and checks that the resulting x -coordinate matches the public key computed by C NaCl:

```

sage: sk=[0x77,0x07,0x6d,0x0a,0x73,0x18,0xa5,0x7d
....:    ,0x3c,0x16,0xc1,0x72,0x51,0xb2,0x66,0x45
....:    ,0xdf,0x4c,0x2f,0x87,0xeb,0xc0,0x99,0x2a
....:    ,0xb1,0x77,0xfb,0xa5,0x1d,0xb9,0x2c,0x2a]
sage: clampsk=sk
sage: clampsk[0]=clampsk[0]-(clampsk[0]%8)
sage: clampsk[31]=64+(clampsk[31]%64)
sage: n=sum(clampsk[i]*256^i for i in range(32))
sage: p=2^255-19
sage: k=GF(p)
sage: E=EllipticCurve([k(0),486662,0,1,0])
sage: s=lift((n*E([k(9),sqrt(k(39420360))]))[0])
sage: pk=[0x85,0x20,0xf0,0x09,0x89,0x30,0xa7,0x54
....:    ,0x74,0x8b,0x7d,0xdc,0xb4,0x3e,0xf7,0x5a
....:    ,0x0d,0xbf,0x3a,0x0d,0x26,0x38,0x1a,0xf4
....:    ,0xeb,0xa4,0xa9,0x8e,0xaa,0x9b,0x4e,0x6a]

```

```
sage: s == sum(pk[i]*256^i for i in range(32))
True
```

Testing: Python vs. scalarmult_curve25519_base. This Python script, contributed by Matthew Dempsky, includes self-contained Curve25519 functions independent of the Sage implementation of elliptic curves:

```
P = 2 ** 255 - 19
A = 486662

def expmod(b, e, m):
    if e == 0: return 1
    t = expmod(b, e / 2, m) ** 2 % m
    if e & 1: t = (t * b) % m
    return t

def inv(x):
    return expmod(x, P - 2, P)

# Addition and doubling formulas taken
# from Appendix D of "Curve25519:
# new Diffie-Hellman speed records".

def add((xn,zn), (xm,zm), (xd,zd)):
    x = 4 * (xm * xn - zm * zn) ** 2 * zd
    z = 4 * (xm * zn - zm * xn) ** 2 * xd
    return (x % P, z % P)

def double((xn,zn)):
    x = (xn ** 2 - zn ** 2) ** 2
    z = 4 * xn * zn * (xn ** 2 + A * xn * zn + zn ** 2)
    return (x % P, z % P)

def curve25519(n, base):
    one = (base,1)
    two = double(one)
    # f(m) evaluates to a tuple
    # containing the mth multiple and the
    # (m+1)th multiple of base.
    def f(m):
        if m == 1: return (one, two)
        (pm, pm1) = f(m / 2)
        if (m & 1):
            return (add(pm, pm1, one), double(pm1))
        return (double(pm), add(pm, pm1, one))
    ((x,z), _) = f(n)
```

```

    return (x * inv(z)) % P

def unpack(s):
    if len(s) != 32:
        raise ValueError('Invalid Curve25519 argument')
    return sum(ord(s[i]) << (8 * i) for i in range(32))

def pack(n):
    return ''.join([chr((n >> (8 * i)) & 255) for i in range(32)])

def clamp(n):
    n &= ~7
    n &= ~(128 << 8 * 31)
    n |= 64 << 8 * 31
    return n

def crypto_scalarmult_curve25519(n, p):
    n = clamp(unpack(n))
    p = unpack(p)
    return pack(curve25519(n, p))

def crypto_scalarmult_curve25519_base(n):
    n = clamp(unpack(n))
    return pack(curve25519(n, 9))

```

After this script the extra commands

```

sk=[0x77,0x07,0x6d,0x0a,0x73,0x18,0xa5,0x7d
    ,0x3c,0x16,0xc1,0x72,0x51,0xb2,0x66,0x45
    ,0xdf,0x4c,0x2f,0x87,0xeb,0xc0,0x99,0x2a
    ,0xb1,0x77,0xfb,0xa5,0x1d,0xb9,0x2c,0x2a]
n=''.join([chr(sk[i]) for i in range(32)])
pk=[0x85,0x20,0xf0,0x09,0x89,0x30,0xa7,0x54
    ,0x74,0x8b,0x7d,0xdc,0xb4,0x3e,0xf7,0x5a
    ,0x0d,0xbf,0x3a,0x0d,0x26,0x38,0x1a,0xf4
    ,0xeb,0xa4,0xa9,0x8e,0xaa,0x9b,0x4e,0x6a]
s=''.join([chr(pk[i]) for i in range(32)])
print s == crypto_scalarmult_curve25519_base(n)

```

print True.

4 Example of the receiver's keys

The following program uses C NaCl to compute the public key corresponding to another secret key:


```

#include <stdio.h>
#include "crypto_scalarmult_curve25519.h"

unsigned char bobsk[32] = {
    0x5d,0xab,0x08,0x7e,0x62,0x4a,0x8a,0x4b
    ,0x79,0xe1,0x7f,0x8b,0x83,0x80,0x0e,0xe6
    ,0x6f,0x3b,0xb1,0x29,0x26,0x18,0xb6,0xfd
    ,0x1c,0x2f,0x8b,0x27,0xff,0x88,0xe0,0xeb
} ;

unsigned char bobpk[32];

main()
{
    int i;
    crypto_scalarmult_curve25519_base(bobpk,bobsk);
    for (i = 0;i < 32;++i) {
        if (i > 0) printf(","); else printf(" ");
        printf("0x%02x",(unsigned int) bobpk[i]);
        if (i % 8 == 7) printf("\n");
    }
    return 0;
}

```

As in the previous section, the secret key bytes embedded into the program were copied from output of `od -t x1 /dev/urandom | head -2`. The output of the program is the corresponding public key:

```

0xde,0x9e,0xdb,0x7d,0x7b,0x7d,0xc1,0xb4
,0xd3,0x5b,0x61,0xc2,0xec,0xe4,0x35,0x37
,0x3f,0x83,0x43,0xc8,0x5b,0x78,0x67,0x4d
,0xad,0xfc,0x7e,0x14,0x6f,0x88,0x2b,0x4f

```

The remaining sections of this document will reuse this example, assuming that Bob's keys are the particular secret key and public key shown here.

Testing: Sage vs. `scalarmult_curve25519_base`. A short Sage script clamps Bob's secret key shown above, converts the result to an integer n , computes $n(9, \sqrt{39420360})$ in $E(\mathbf{F}_p)$, and checks that the resulting x -coordinate matches the public key computed by C NaCl:

```

sage: sk=[0x5d,0xab,0x08,0x7e,0x62,0x4a,0x8a,0x4b
....:    ,0x79,0xe1,0x7f,0x8b,0x83,0x80,0x0e,0xe6
....:    ,0x6f,0x3b,0xb1,0x29,0x26,0x18,0xb6,0xfd
....:    ,0x1c,0x2f,0x8b,0x27,0xff,0x88,0xe0,0xeb]
sage: clampsk=sk
sage: clampsk[0]=clampsk[0]-(clampsk[0]%8)
sage: clampsk[31]=64+(clampsk[31]%64)

```

```

sage: n=sum(clampsk[i]*256^i for i in range(32))
sage: p=2^255-19
sage: k=GF(p)
sage: E=EllipticCurve([k(0),486662,0,1,0])
sage: s=lift((n*E([k(9),sqrt(k(39420360))]))[0])
sage: pk=[0xde,0x9e,0xdb,0x7d,0x7b,0x7d,0xc1,0xb4
.....:      ,0xd3,0x5b,0x61,0xc2,0xec,0xe4,0x35,0x37
.....:      ,0x3f,0x83,0x43,0xc8,0x5b,0x78,0x67,0x4d
.....:      ,0xad,0xfc,0x7e,0x14,0x6f,0x88,0x2b,0x4f]
sage: s == sum(pk[i]*256^i for i in range(32))
True

```

Testing: Python vs. scalarmult_curve25519_base. After the Python script shown in Section 3, the extra commands

```

sk=[0x5d,0xab,0x08,0x7e,0x62,0x4a,0x8a,0x4b
    ,0x79,0xe1,0x7f,0x8b,0x83,0x80,0x0e,0xe6
    ,0x6f,0x3b,0xb1,0x29,0x26,0x18,0xb6,0xfd
    ,0x1c,0x2f,0x8b,0x27,0xff,0x88,0xe0,0xeb]
n='' .join([chr(sk[i]) for i in range(32)])
pk=[0xde,0x9e,0xdb,0x7d,0x7b,0x7d,0xc1,0xb4
    ,0xd3,0x5b,0x61,0xc2,0xec,0xe4,0x35,0x37
    ,0x3f,0x83,0x43,0xc8,0x5b,0x78,0x67,0x4d
    ,0xad,0xfc,0x7e,0x14,0x6f,0x88,0x2b,0x4f]
s='' .join([chr(pk[i]) for i in range(32)])
print s == crypto_scalarmult_curve25519_base(n)

```

print True.

5 Shared secret

At this point Alice has a secret key $a \in \{0, 1, \dots, 255\}^{32}$ and a public key $A = \text{Curve25519}(\text{ClampC}(a), \underline{9}) \in \{0, 1, \dots, 255\}^{32}$. Similarly, Bob has a secret key b and a public key $B = \text{Curve25519}(\text{ClampC}(b), \underline{9})$.

Assume that Alice knows Bob's public key from a previous secure channel—for example, from meeting privately with Bob. Similarly assume that Bob knows Alice's public key. There is no hope of security if the previous channel allows forgeries: for example, if an attacker can replace Bob's public key with the attacker's public key then Alice will end up encrypting a packet to the attacker instead of to Bob.

Alice computes $\text{Curve25519}(\text{ClampC}(a), B)$ from her secret key a and Bob's public key B . Bob computes $\text{Curve25519}(\text{ClampC}(b), A)$ from his secret key b and Alice's public key A . The definition of Curve25519 immediately implies that $\text{Curve25519}(\text{ClampC}(a), B) = \text{Curve25519}(\text{ClampC}(b), A)$, so at this point Alice and Bob have computed the same 32-byte string.

In the next step, described in Section 7, Alice will convert this 32-byte shared secret k into a 32-byte string $\text{HSalsa20}(k, 0)$, which is then used to encrypt and authenticate packets. Bob similarly uses $\text{HSalsa20}(k, 0)$ to verify and decrypt the packets. No other use is made of k . One can thus view $\text{HSalsa20}(k, 0)$ as the shared secret rather than k .

Security notes beyond ECDLP. An attacker who can solve the elliptic-curve discrete-logarithm problem can figure out Alice’s secret key from Alice’s public key, and can then compute the shared secret the same way Alice does; or figure out Bob’s secret key from Bob’s public key, and can then compute the shared secret the same way Bob does.

Computing the shared secret from the two public keys—the “Diffie–Hellman problem”—is widely conjectured to be as difficult as computing discrete logarithms. There are weak theorems along these lines, stating that (for typical elliptic curves) a reliable algorithm to solve the Diffie–Hellman problem can be converted into a discrete-logarithm algorithm costing about ten thousand times as much.

It is much easier to compute *some information* about the 32-byte string k . There are only $p_1 \approx 2^{251}$ possibilities for k , and the set of possibilities for k is an easy-to-recognize set: for example, the last bit of k is always 0. However, $\text{HSalsa20}(k, 0)$ is conjectured to be indistinguishable from $\text{HSalsa20}(k', 0)$ where k' is a uniform random Curve25519 output.

It is often conjectured that the “decision Diffie–Hellman problem” is hard: i.e., that k is indistinguishable from k' . However, this DDH conjecture is overkill. What matters is that $\text{HSalsa20}(k, 0)$ is indistinguishable from $\text{HSalsa20}(k', 0)$.

Alice can reuse her secret key and public key for communicating with many parties. Some of those parties may be attackers with fake public keys—32-byte strings that are not of the form $\text{Curve25519}(\text{ClampC}(\dots), \underline{9})$. The corresponding points can be in the “twist group” $E(\mathbf{F}_{p^2}) \cap (\{\infty\} \cup (\mathbf{F}_p \times \sqrt{2}\mathbf{F}_p))$; even if the points are in $E(\mathbf{F}_p)$, they can be outside the subgroup of order p_1 . If the points have small order then they can reveal Alice’s secret n modulo that order. Fortunately, $E(\mathbf{F}_p)$ has order $8p_1$ by the Hasse–Weil theorem, and the twist group has order $4p_2$ where p_2 is the prime number

$$2^{253} - 55484635554744707071703875581767296995 = (p + 1)/2 - 2p_1.$$

The following Sage transcript captures the relevant facts about p_2 :

```
sage: p=2^255-19
sage: p1=2^252+2774231777372353535851937790883648493
sage: p2=2^253-55484635554744707071703875581767296995
sage: p2.is_prime()
True
sage: 8*p1+4*p2-2*(p+1)
0
```

Consequently the only possible small orders are 1, 2, 4, and 8, and an attacker can learn at most Alice’s $n \bmod 8$, which is always 0 by construction. See [8, Section 3] for further discussion of active attacks and twist security.

6 Example of the shared secret

The following program, starting from Section 3's example of Alice's secret key and Section 4's example of Bob's public key, uses C NaCl to compute the secret shared between Alice and Bob:

```
#include <stdio.h>
#include "crypto_scalarmult_curve25519.h"

unsigned char alicesk[32] = {
    0x77,0x07,0x6d,0x0a,0x73,0x18,0xa5,0x7d
    ,0x3c,0x16,0xc1,0x72,0x51,0xb2,0x66,0x45
    ,0xdf,0x4c,0x2f,0x87,0xeb,0xc0,0x99,0x2a
    ,0xb1,0x77,0xfb,0xa5,0x1d,0xb9,0x2c,0x2a
} ;

unsigned char bobpk[32] = {
    0xde,0x9e,0xdb,0x7d,0x7b,0x7d,0xc1,0xb4
    ,0xd3,0x5b,0x61,0xc2,0xec,0xe4,0x35,0x37
    ,0x3f,0x83,0x43,0xc8,0x5b,0x78,0x67,0x4d
    ,0xad,0xfc,0x7e,0x14,0x6f,0x88,0x2b,0x4f
} ;

unsigned char k[32];

main()
{
    int i;
    crypto_scalarmult_curve25519(k,alicesk,bobpk);
    for (i = 0;i < 32;++i) {
        if (i > 0) printf(","); else printf(" ");
        printf("0x%02x",(unsigned int) k[i]);
        if (i % 8 == 7) printf("\n");
    }
    return 0;
}
```

The program produces the following output:

```
0x4a,0x5d,0x9d,0x5b,0xa4,0xce,0x2d,0xe1
,0x72,0x8e,0x3b,0xf4,0x80,0x35,0x0f,0x25
,0xe0,0x7e,0x21,0xc9,0x47,0xd1,0x9e,0x33
,0x76,0xf0,0x9b,0x3c,0x1e,0x16,0x17,0x42
```

The following program, starting from Section 4's example of Bob's secret key and Section 3's example of Alice's public key, uses C NaCl to compute the secret shared between Alice and Bob:

```

#include <stdio.h>
#include "crypto_scalarmult_curve25519.h"

unsigned char bobsk[32] = {
    0x5d,0xab,0x08,0x7e,0x62,0x4a,0x8a,0x4b
    ,0x79,0xe1,0x7f,0x8b,0x83,0x80,0x0e,0xe6
    ,0x6f,0x3b,0xb1,0x29,0x26,0x18,0xb6,0xfd
    ,0x1c,0x2f,0x8b,0x27,0xff,0x88,0xe0,0xeb
} ;

unsigned char alicepk[32] = {
    0x85,0x20,0xf0,0x09,0x89,0x30,0xa7,0x54
    ,0x74,0x8b,0x7d,0xdc,0xb4,0x3e,0xf7,0x5a
    ,0x0d,0xbf,0x3a,0x0d,0x26,0x38,0x1a,0xf4
    ,0xeb,0xa4,0xa9,0x8e,0xaa,0x9b,0x4e,0x6a
} ;

unsigned char k[32];

main()
{
    int i;
    crypto_scalarmult_curve25519(k,bobsk,alicepk);
    for (i = 0;i < 32;++i) {
        if (i > 0) printf(","); else printf(" ");
        printf("0x%02x",(unsigned int) k[i]);
        if (i % 8 == 7) printf("\n");
    }
    return 0;
}

```

This program produces the same output as the previous program.

Testing: Sage vs. scalarmult_curve25519. A short Sage script clamps Alice's secret key, converts the result to an integer n , clamps Bob's secret key, converts the result to an integer m , computes $mn(9, \sqrt{39420360})$ in $E(\mathbf{F}_p)$, and checks that the x -coordinate of the result matches the shared secret computed by C NaCl:

```

sage: alicesk=[0x77,0x07,0x6d,0x0a,0x73,0x18,0xa5,0x7d
....:         ,0x3c,0x16,0xc1,0x72,0x51,0xb2,0x66,0x45
....:         ,0xdf,0x4c,0x2f,0x87,0xeb,0xc0,0x99,0x2a
....:         ,0xb1,0x77,0xfb,0xa5,0x1d,0xb9,0x2c,0x2a]
sage: clampsk=alicesk
sage: clampsk[0]=clampsk[0]-(clampsk[0]%8)
sage: clampsk[31]=64+(clampsk[31]%64)
sage: n=sum(clampsk[i]*256^i for i in range(32))

```

```

sage: bobsk=[0x5d,0xab,0x08,0x7e,0x62,0x4a,0x8a,0x4b
....:      ,0x79,0xe1,0x7f,0x8b,0x83,0x80,0x0e,0xe6
....:      ,0x6f,0x3b,0xb1,0x29,0x26,0x18,0xb6,0xfd
....:      ,0x1c,0x2f,0x8b,0x27,0xff,0x88,0xe0,0xeb]
sage: clampsk=bobsk
sage: clampsk[0]=clampsk[0]-(clampsk[0]%8)
sage: clampsk[31]=64+(clampsk[31]%64)
sage: m=sum(clampsk[i]*256^i for i in range(32))
sage: p=2^255-19
sage: k=GF(p)
sage: E=EllipticCurve([k(0),486662,0,1,0])
sage: s=lift((m*n*E([k(9),sqrt(k(39420360))]))[0])
sage: shared=[0x4a,0x5d,0x9d,0x5b,0xa4,0xce,0x2d,0xe1
....:      ,0x72,0x8e,0x3b,0xf4,0x80,0x35,0x0f,0x25
....:      ,0xe0,0x7e,0x21,0xc9,0x47,0xd1,0x9e,0x33
....:      ,0x76,0xf0,0x9b,0x3c,0x1e,0x16,0x17,0x42]
sage: s == sum(shared[i]*256^i for i in range(32))
True

```

Testing: Python vs. scalarmult_curve25519. After the Python script shown in Section 3, the extra commands

```

alicesk=[0x77,0x07,0x6d,0x0a,0x73,0x18,0xa5,0x7d
....:      ,0x3c,0x16,0xc1,0x72,0x51,0xb2,0x66,0x45
....:      ,0xdf,0x4c,0x2f,0x87,0xeb,0xc0,0x99,0x2a
....:      ,0xb1,0x77,0xfb,0xa5,0x1d,0xb9,0x2c,0x2a]
a='' .join([chr(alicesk[i]) for i in range(32)])
bobpk=[0xde,0x9e,0xdb,0x7d,0x7b,0x7d,0xc1,0xb4
....:      ,0xd3,0x5b,0x61,0xc2,0xec,0xe4,0x35,0x37
....:      ,0x3f,0x83,0x43,0xc8,0x5b,0x78,0x67,0x4d
....:      ,0xad,0xfc,0x7e,0x14,0x6f,0x88,0x2b,0x4f]
b='' .join([chr(bobpk[i]) for i in range(32)])
shared=[0x4a,0x5d,0x9d,0x5b,0xa4,0xce,0x2d,0xe1
....:      ,0x72,0x8e,0x3b,0xf4,0x80,0x35,0x0f,0x25
....:      ,0xe0,0x7e,0x21,0xc9,0x47,0xd1,0x9e,0x33
....:      ,0x76,0xf0,0x9b,0x3c,0x1e,0x16,0x17,0x42]
s='' .join([chr(shared[i]) for i in range(32)])
print s == crypto_scalarmult_curve25519(a,b)

```

print true, and the extra commands

```

bobsk=[0x5d,0xab,0x08,0x7e,0x62,0x4a,0x8a,0x4b
....:      ,0x79,0xe1,0x7f,0x8b,0x83,0x80,0x0e,0xe6
....:      ,0x6f,0x3b,0xb1,0x29,0x26,0x18,0xb6,0xfd
....:      ,0x1c,0x2f,0x8b,0x27,0xff,0x88,0xe0,0xeb]
b='' .join([chr(bobsk[i]) for i in range(32)])
alicepk=[0x85,0x20,0xf0,0x09,0x89,0x30,0xa7,0x54

```

```

        ,0x74,0x8b,0x7d,0xdc,0xb4,0x3e,0xf7,0x5a
        ,0x0d,0xbf,0x3a,0x0d,0x26,0x38,0x1a,0xf4
        ,0xeb,0xa4,0xa9,0x8e,0xaa,0x9b,0x4e,0x6a]
a='' .join([chr(alicepk[i]) for i in range(32)])
shared=[0x4a,0x5d,0x9d,0x5b,0xa4,0xce,0x2d,0xe1
        ,0x72,0x8e,0x3b,0xf4,0x80,0x35,0x0f,0x25
        ,0xe0,0x7e,0x21,0xc9,0x47,0xd1,0x9e,0x33
        ,0x76,0xf0,0x9b,0x3c,0x1e,0x16,0x17,0x42]
s='' .join([chr(shared[i]) for i in range(32)])
print s == crypto_scalarmult_curve25519(b,a)

```

print true.

7 Nonce and stream

At this point Alice and Bob have a shared secret $k \in \{0, 1, \dots, 255\}^{32}$. This secret can be used to protect a practically infinite sequence of packets exchanged between Alice and Bob.

Alice and Bob assign to each packet a **nonce** $n \in \{0, 1, \dots, 255\}^{24}$: a unique message number that will never be reused for other packets exchanged between Alice and Bob. For example, the nonce can be chosen as a simple counter: 0 for Alice’s first packet, 1 for Bob’s first packet, 2 for Alice’s second packet, 3 for Bob’s second packet, 4 for Alice’s third packet, 5 for Bob’s third packet, etc. Choosing the nonce as a counter followed by (e.g.) 32 random bits helps protect some protocols against denial-of-service attacks. In many applications it is better to increase the counter to, e.g., the number of nanoseconds that have passed since a standard epoch in the local clock, so that the current value of the counter does not leak the traffic rate. Note that “increase” does not mean “increase or decrease”; if the clock jumps backwards, the counter must continue to increase.

Alice uses the shared secret k to expand the nonce n into a long stream. Specifically, Alice computes a first-level key $\text{HSalsa20}(k, 0)$; uses the first 16 bytes n_1 of the nonce to compute a second-level key $\text{HSalsa20}(\text{HSalsa20}(k, 0), n_1)$; and uses the remaining 8 bytes n_2 of the nonce to compute a long stream $\text{Salsa20}(\text{HSalsa20}(\text{HSalsa20}(k, 0), n_1), n_2)$. This stream is then used to encrypt and authenticate the packet, as described in subsequent sections.

This section defines HSalsa20 and Salsa20. Many of the definitions here are copied from the original Salsa20 specification [7]. Section 8 gives an example of nonce expansion, starting from the key examples used in Sections 4, 3, and 6.

Words. A **word** is an element of $\{0, 1, \dots, 2^{32} - 1\}$.

The **sum** of two words u, v is $u + v \bmod 2^{32}$. The sum is denoted $u + v$; there is no risk of confusion. For example, $0xc0a8787e + 0x9fd1161d = 0x60798e9b$.

The **exclusive-or** of two words u, v , denoted $u \oplus v$, is the sum of u and v with carries suppressed. In other words, if $u = \sum_i 2^i u_i$ and $v = \sum_i 2^i v_i$ then $u \oplus v = \sum_i 2^i (u_i + v_i - 2u_i v_i)$. For example, $0xc0a8787e \oplus 0x9fd1161d = 0x5f796e63$.

For each $c \in \{0, 1, 2, 3, \dots\}$, the c -bit left rotation of a word u , denoted $u \lll c$, is the unique nonzero word congruent to $2^c u$ modulo $2^{32} - 1$, except that $0 \lll c = 0$. In other words, if $u = \sum_i 2^i u_i$ then $u \lll c = \sum_i 2^{i+c \bmod 32} u_i$. For example, `0xc0a8787e` $\lll 5 =$ `0x150f0fd8`.

The quarterround function. If $y = (y_0, y_1, y_2, y_3) \in \{0, 1, \dots, 2^{32} - 1\}^4$ then $\text{quarterround}(y) \in \{0, 1, \dots, 2^{32} - 1\}^4$ is defined as (z_0, z_1, z_2, z_3) where

$$\begin{aligned} z_1 &= y_1 \oplus ((y_0 + y_3) \lll 7), \\ z_2 &= y_2 \oplus ((z_1 + y_0) \lll 9), \\ z_3 &= y_3 \oplus ((z_2 + z_1) \lll 13), \\ z_0 &= y_0 \oplus ((z_3 + z_2) \lll 18). \end{aligned}$$

The rowround function. If $y = (y_0, y_1, y_2, y_3, \dots, y_{15}) \in \{0, 1, \dots, 2^{32} - 1\}^{16}$ then $\text{rowround}(y) \in \{0, 1, \dots, 2^{32} - 1\}^{16}$ is defined as $(z_0, z_1, z_2, z_3, \dots, z_{15})$ where

$$\begin{aligned} (z_0, z_1, z_2, z_3) &= \text{quarterround}(y_0, y_1, y_2, y_3), \\ (z_5, z_6, z_7, z_4) &= \text{quarterround}(y_5, y_6, y_7, y_4), \\ (z_{10}, z_{11}, z_8, z_9) &= \text{quarterround}(y_{10}, y_{11}, y_8, y_9), \\ (z_{15}, z_{12}, z_{13}, z_{14}) &= \text{quarterround}(y_{15}, y_{12}, y_{13}, y_{14}). \end{aligned}$$

The columnround function. If $x = (x_0, x_1, \dots, x_{15}) \in \{0, 1, \dots, 2^{32} - 1\}^{16}$ then $\text{columnround}(x) \in \{0, 1, \dots, 2^{32} - 1\}^{16}$ is defined as $(y_0, y_1, y_2, y_3, \dots, y_{15})$ where

$$\begin{aligned} (y_0, y_4, y_8, y_{12}) &= \text{quarterround}(x_0, x_4, x_8, x_{12}), \\ (y_5, y_9, y_{13}, y_1) &= \text{quarterround}(x_5, x_9, x_{13}, x_1), \\ (y_{10}, y_{14}, y_2, y_6) &= \text{quarterround}(x_{10}, x_{14}, x_2, x_6), \\ (y_{15}, y_3, y_7, y_{11}) &= \text{quarterround}(x_{15}, x_3, x_7, x_{11}). \end{aligned}$$

Equivalent formula: $(y_0, y_4, y_8, y_{12}, y_1, y_5, y_9, y_{13}, y_2, y_6, y_{10}, y_{14}, y_3, y_7, y_{11}, y_{15}) = \text{rowround}(x_0, x_4, x_8, x_{12}, x_1, x_5, x_9, x_{13}, x_2, x_6, x_{10}, x_{14}, x_3, x_7, x_{11}, x_{15})$.

The doubleround function. If $x \in \{0, 1, \dots, 2^{32} - 1\}^{16}$ then $\text{doubleround}(x) \in \{0, 1, \dots, 2^{32} - 1\}^{16}$ is defined as $\text{rowround}(\text{columnround}(x))$.

The littleendian function. If $b = (b_0, b_1, b_2, b_3) \in \{0, 1, 2, 3, \dots, 255\}^4$ then $\text{littleendian}(b) \in \{0, 1, \dots, 2^{32} - 1\}$ is defined as $b_0 + 2^8 b_1 + 2^{16} b_2 + 2^{24} b_3$. More generally, if $b = (b_0, b_1, \dots, b_{4k-1}) \in \{0, 1, \dots, 255\}^{4k}$ then $\text{littleendian}(b) \in \{0, 1, \dots, 2^{32} - 1\}^k$ is defined as

$$(b_0 + 2^8 b_1 + 2^{16} b_2 + 2^{24} b_3, b_4 + 2^8 b_5 + 2^{16} b_6 + 2^{24} b_7, \dots).$$

Note that littleendian is invertible.

The HSalsa20 function. The function

$$\text{HSalsa20} : \{0, 1, \dots, 255\}^{32} \times \{0, 1, \dots, 255\}^{16} \rightarrow \{0, 1, \dots, 255\}^{32}$$

is defined as follows.

Fix $k \in \{0, 1, \dots, 255\}^{32}$ and $n \in \{0, 1, \dots, 255\}^{16}$. Define $(x_0, x_1, \dots, x_{15}) \in \{0, 1, \dots, 2^{32} - 1\}^{16}$ as follows:

- $(x_0, x_5, x_{10}, x_{15}) = (\text{0x61707865}, \text{0x3320646e}, \text{0x79622d32}, \text{0x6b206574})$; in other words, $(x_0, x_5, x_{10}, x_{15})$ is the **Salsa20 constant**.
- $(x_1, x_2, x_3, x_4, x_{11}, x_{12}, x_{13}, x_{14}) = \text{littleendian}(k)$; and
- $(x_6, x_7, x_8, x_9) = \text{littleendian}(n)$.

Define $(z_0, z_1, \dots, z_{15}) = \text{doubleround}^{10}(x_0, x_1, \dots, x_{15})$. Then $\text{HSalsa20}(k, n) = \text{littleendian}^{-1}(z_0, z_5, z_{10}, z_{15}, z_6, z_7, z_8, z_9)$.

The Salsa20 expansion function. The function

$$\text{Salsa20} : \{0, 1, \dots, 255\}^{32} \times \{0, 1, \dots, 255\}^{16} \rightarrow \{0, 1, \dots, 255\}^{64}$$

is defined as follows.

Fix $k \in \{0, 1, \dots, 255\}^{32}$ and $n \in \{0, 1, \dots, 255\}^{16}$. Define $(x_0, x_1, \dots, x_{15}) \in \{0, 1, \dots, 2^{32} - 1\}^{16}$ as follows:

- $(x_0, x_5, x_{10}, x_{15})$ is the Salsa20 constant;
- $(x_1, x_2, x_3, x_4, x_{11}, x_{12}, x_{13}, x_{14}) = \text{littleendian}(k)$; and
- $(x_6, x_7, x_8, x_9) = \text{littleendian}(n)$.

Define $(z_0, z_1, \dots, z_{15}) = \text{doubleround}^{10}(x_0, x_1, \dots, x_{15})$. Then $\text{Salsa20}(k, n) = \text{littleendian}^{-1}(x_0 + z_0, x_1 + z_1, \dots, x_{15} + z_{15})$.

The Salsa20 streaming function. The function

$$\text{Salsa20} : \{0, 1, \dots, 255\}^{32} \times \{0, 1, \dots, 255\}^8 \rightarrow \{0, 1, \dots, 255\}^{2^{70}}$$

is defined as follows: $\text{Salsa20}(k, n) = \text{Salsa20}(k, n, \underline{0}), \text{Salsa20}(k, n, \underline{1}), \dots$. Here \underline{b} means the 8-byte string $(b \bmod 256, \lfloor b/256 \rfloor \bmod 256, \dots)$.

Security notes. ECRYPT, a consortium of European research organizations, issued a Call for Stream Cipher Primitives in November 2004, and received 34 submissions from 97 cryptographers in 19 countries. In April 2008, after two hundred papers and several conferences, ECRYPT selected a portfolio of 4 software ciphers and 4 lower-security hardware ciphers.

I submitted Salsa20. Later I suggested the reduced-round variants Salsa20/12 and Salsa20/8 (replacing doubleround^{10} with doubleround^6 and doubleround^4 respectively) as higher-speed options for users who value speed more highly than confidence. Four attack papers by fourteen cryptanalysts ([12], [13], [21], and [3]) culminated in a 2^{184} -operation attack on Salsa20/7 and a 2^{251} -operation attack on Salsa20/8. The eSTREAM portfolio recommended Salsa20/12: “Eight and twenty round versions were also considered during the eSTREAM process, but we

feel that Salsa20/12 offers the best balance, combining a very nice performance profile with what appears to be a comfortable margin for security.”

The standard (“PRF”) security conjecture for Salsa20 is that the Salsa20 output blocks, for a uniform random secret key k , are indistinguishable from independent uniform random 64-byte strings. This conjecture implies the analogous security conjecture for HSalsa20: by [10, Theorem 3.3], any attack against HSalsa20 can be converted into an attack against Salsa20 having exactly the same effectiveness and essentially the same speed.

This conjecture also implies an analogous security conjecture for the cascade $(n_1, n_2) \mapsto \text{Salsa20}(\text{HSalsa20}(\text{HSalsa20}(k, 0), n_1), n_2)$: by [10, Theorem 3.1], any q -query attack against the cascade can be converted into an attack against Salsa20 having at least $1/(2q + 1)$ as much effectiveness and essentially the same speed.

A Curve25519 output k is not a uniform random 32-byte string, but any attack against a uniform random Curve25519 output can be converted into an attack against a uniform random 32-byte string having at least $1/32$ as much effectiveness and essentially the same speed—and therefore an attack against Salsa20 having at least $1/(64q + 32)$ as much effectiveness and essentially the same speed.

8 Example of the long stream

The following program starts from Section 6’s example of a shared secret k and uses C NaCl to compute the first-level key $k_1 = \text{HSalsa20}(k, 0)$:

```
#include <stdio.h>
#include "crypto_core_hsalsa20.h"

unsigned char shared[32] = {
    0x4a,0x5d,0x9d,0x5b,0xa4,0xce,0x2d,0xe1
    ,0x72,0x8e,0x3b,0xf4,0x80,0x35,0x0f,0x25
    ,0xe0,0x7e,0x21,0xc9,0x47,0xd1,0x9e,0x33
    ,0x76,0xf0,0x9b,0x3c,0x1e,0x16,0x17,0x42
} ;

unsigned char zero[32] = { 0 };

unsigned char c[16] = {
    0x65,0x78,0x70,0x61,0x6e,0x64,0x20,0x33
    ,0x32,0x2d,0x62,0x79,0x74,0x65,0x20,0x6b
} ;

unsigned char firstkey[32];

main()
{
```

```

int i;
crypto_core_hsalsa20(firstkey,zero,shared,c);
for (i = 0;i < 32;++i) {
    if (i > 0) printf(","); else printf(" ");
    printf("0x%02x",(unsigned int) firstkey[i]);
    if (i % 8 == 7) printf("\n");
}
return 0;
}

```

The program prints the following output:

```

0x1b,0x27,0x55,0x64,0x73,0xe9,0x85,0xd4
,0x62,0xcd,0x51,0x19,0x7a,0x9a,0x46,0xc7
,0x60,0x09,0x54,0x9e,0xac,0x64,0x74,0xf2
,0x06,0xc4,0xee,0x08,0x44,0xf6,0x83,0x89

```

The following program starts from this k_1 example and a sample nonce prefix n_1 , and uses C NaCl to compute the second-level key $k_2 = \text{HSalsa20}(k_1, n_1)$:

```

#include <stdio.h>
#include "crypto_core_hsalsa20.h"

unsigned char firstkey[32] = {
    0x1b,0x27,0x55,0x64,0x73,0xe9,0x85,0xd4
,0x62,0xcd,0x51,0x19,0x7a,0x9a,0x46,0xc7
,0x60,0x09,0x54,0x9e,0xac,0x64,0x74,0xf2
,0x06,0xc4,0xee,0x08,0x44,0xf6,0x83,0x89
} ;

unsigned char nonceprefix[16] = {
    0x69,0x69,0x6e,0xe9,0x55,0xb6,0x2b,0x73
,0xcd,0x62,0xbd,0xa8,0x75,0xfc,0x73,0xd6
} ;

unsigned char c[16] = {
    0x65,0x78,0x70,0x61,0x6e,0x64,0x20,0x33
,0x32,0x2d,0x62,0x79,0x74,0x65,0x20,0x6b
} ;

unsigned char secondkey[32];

main()
{
    int i;
    crypto_core_hsalsa20(secondkey,nonceprefix,firstkey,c);
    for (i = 0;i < 32;++i) {

```

```

        if (i > 0) printf(","); else printf(" ");
        printf("0x%02x", (unsigned int) secondkey[i]);
        if (i % 8 == 7) printf("\n");
    }
    return 0;
}

```

The program prints the following output:

```

0xdc,0x90,0x8d,0xda,0x0b,0x93,0x44,0xa9
,0x53,0x62,0x9b,0x73,0x38,0x20,0x77,0x88
,0x80,0xf3,0xce,0xb4,0x21,0xbb,0x61,0xb9
,0x1c,0xbd,0x4c,0x3e,0x66,0x25,0x6c,0xe4

```

The following program starts from this k_2 example and an example of a nonce suffix n_2 , and uses C NaCl to print (in binary format) the first 4194304 bytes of $\text{Salsa20}(k_2, n_2)$:

```

#include <stdio.h>
#include "crypto_core_salsa20.h"

unsigned char secondkey[32] = {
    0xdc,0x90,0x8d,0xda,0x0b,0x93,0x44,0xa9
,0x53,0x62,0x9b,0x73,0x38,0x20,0x77,0x88
,0x80,0xf3,0xce,0xb4,0x21,0xbb,0x61,0xb9
,0x1c,0xbd,0x4c,0x3e,0x66,0x25,0x6c,0xe4
} ;

unsigned char noncesuffix[8] = {
    0x82,0x19,0xe0,0x03,0x6b,0x7a,0x0b,0x37
} ;

unsigned char c[16] = {
    0x65,0x78,0x70,0x61,0x6e,0x64,0x20,0x33
,0x32,0x2d,0x62,0x79,0x74,0x65,0x20,0x6b
} ;

unsigned char in[16] = { 0 } ;

unsigned char outputblock[64];

main()
{
    int i;
    for (i = 0; i < 8; ++i) in[i] = noncesuffix[i];
    do {
        do {

```

```

        crypto_core_salsa20(outputblock,in,secondkey,c);
        for (i = 0;i < 64;++i) putchar(outputblock[i]);
    } while (++in[8]);
} while (++in[9]);
return 0;
}

```

662b9d0e3463029156069b12f918691a98f7dfb2ca0393c96bbfc6b1fbd630a2 is the SHA-256 checksum of the output.

Testing: `core_salsa20` vs. `stream_salsa20`. The following program has the same output as the previous program, but uses `crypto_stream_salsa20` to generate the entire output stream at once:

```

#include <stdio.h>
#include "crypto_stream_salsa20.h"

unsigned char secondkey[32] = {
    0xdc,0x90,0x8d,0xda,0x0b,0x93,0x44,0xa9
    ,0x53,0x62,0x9b,0x73,0x38,0x20,0x77,0x88
    ,0x80,0xf3,0xce,0xb4,0x21,0xbb,0x61,0xb9
    ,0x1c,0xbd,0x4c,0x3e,0x66,0x25,0x6c,0xe4
} ;

unsigned char noncesuffix[8] = {
    0x82,0x19,0xe0,0x03,0x6b,0x7a,0x0b,0x37
} ;

unsigned char output[4194304];

main()
{
    int i;
    crypto_stream_salsa20(output,4194304,noncesuffix,secondkey);
    for (i = 0;i < 4194304;++i) putchar(output[i]);
    return 0;
}

```

Testing: `core_salsa20` vs. `stream_xsalsa20`. The following program has the same output as the previous two programs, but uses `crypto_stream_xsalsa20` to generate the entire output stream starting from the first-level key k_1 and the complete nonce $n = (n_1, n_2)$:

```

#include <stdio.h>
#include "crypto_stream_xsalsa20.h"

unsigned char firstkey[32] = {

```

```

    0x1b,0x27,0x55,0x64,0x73,0xe9,0x85,0xd4
    ,0x62,0xcd,0x51,0x19,0x7a,0x9a,0x46,0xc7
    ,0x60,0x09,0x54,0x9e,0xac,0x64,0x74,0xf2
    ,0x06,0xc4,0xee,0x08,0x44,0xf6,0x83,0x89
} ;

unsigned char nonce[24] = {
    0x69,0x69,0x6e,0xe9,0x55,0xb6,0x2b,0x73
    ,0xcd,0x62,0xbd,0xa8,0x75,0xfc,0x73,0xd6
    ,0x82,0x19,0xe0,0x03,0x6b,0x7a,0x0b,0x37
} ;

unsigned char output[4194304];

main()
{
    int i;
    crypto_stream_xsalsa20(output,4194304,nonce,firstkey);
    for (i = 0;i < 4194304;++i) putchar(output[i]);
    return 0;
}

```

Testing: Python vs. core_hsalsa20. The following Python script, based in part on a script contributed by Matthew Dempsky, computes HSalsa20(k , 0) and compares the result to the k_1 computed by C NaCl:

```

import struct

def rotate(x, n):
    x &= 0xffffffff
    return ((x << n) | (x >> (32 - n))) & 0xffffffff

def step(s, i, j, k, r):
    s[i] ^= rotate(s[j] + s[k],r)

def quarterround(s, i0, i1, i2, i3):
    step(s, i1, i0, i3, 7)
    step(s, i2, i1, i0, 9)
    step(s, i3, i2, i1, 13)
    step(s, i0, i3, i2, 18)

def rowround(s):
    quarterround(s, 0, 1, 2, 3)
    quarterround(s, 5, 6, 7, 4)
    quarterround(s, 10, 11, 8, 9)
    quarterround(s, 15, 12, 13, 14)

```

```

def columnround(s):
    quarterround(s, 0, 4, 8, 12)
    quarterround(s, 5, 9, 13, 1)
    quarterround(s, 10, 14, 2, 6)
    quarterround(s, 15, 3, 7, 11)

def doubleround(s):
    columnround(s)
    rowround(s)

def hsalsa20(n,k):
    n=''.join([chr(n[i]) for i in range(16)])
    n = struct.unpack('<4I', n)
    k=''.join([chr(k[i]) for i in range(32)])
    k = struct.unpack('<8I', k)
    s = [0] * 16
    s[:5] = struct.unpack('<4I', 'expand 32-byte k')
    s[1:5] = k[:4]
    s[6:10] = n
    s[11:15] = k[4:]
    for i in range(10): doubleround(s)
    s = [s[i] for i in [0,5,10,15,6,7,8,9]]
    return struct.pack('<8I',*s)

k = [0x4a,0x5d,0x9d,0x5b,0xa4,0xce,0x2d,0xe1
     ,0x72,0x8e,0x3b,0xf4,0x80,0x35,0x0f,0x25
     ,0xe0,0x7e,0x21,0xc9,0x47,0xd1,0x9e,0x33
     ,0x76,0xf0,0x9b,0x3c,0x1e,0x16,0x17,0x42]
n = [0] * 16

expected=[0x1b,0x27,0x55,0x64,0x73,0xe9,0x85,0xd4
         ,0x62,0xcd,0x51,0x19,0x7a,0x9a,0x46,0xc7
         ,0x60,0x09,0x54,0x9e,0xac,0x64,0x74,0xf2
         ,0x06,0xc4,0xee,0x08,0x44,0xf6,0x83,0x89]
expected = ''.join([chr(expected[i]) for i in range(32)])

print hsalsa20(n,k) == expected

```

The script prints True.

The following extra commands compute HSalsa20(k_1, n_1), where n_1 is the nonce prefix shown above, and compare the result to the k_2 computed by C NaCl:

```

k=[0x1b,0x27,0x55,0x64,0x73,0xe9,0x85,0xd4
  ,0x62,0xcd,0x51,0x19,0x7a,0x9a,0x46,0xc7
  ,0x60,0x09,0x54,0x9e,0xac,0x64,0x74,0xf2

```

```

    ,0x06,0xc4,0xee,0x08,0x44,0xf6,0x83,0x89]
n=[0x69,0x69,0x6e,0xe9,0x55,0xb6,0x2b,0x73
   ,0xcd,0x62,0xbd,0xa8,0x75,0xfc,0x73,0xd6]

expected = [0xdc,0x90,0x8d,0xda,0x0b,0x93,0x44,0xa9
            ,0x53,0x62,0x9b,0x73,0x38,0x20,0x77,0x88
            ,0x80,0xf3,0xce,0xb4,0x21,0xbb,0x61,0xb9
            ,0x1c,0xbd,0x4c,0x3e,0x66,0x25,0x6c,0xe4]
expected = ''.join([chr(expected[i]) for i in range(32)])

print hsalsa20(n,k) == expected

```

These commands print True.

Testing: Python vs. stream_salsa20. The following Python script, based in part on a script contributed by Matthew Dempsky, computes the first 4194304 bytes of Salsa20(k_2, n_2), for the sample k_2, n_2 shown above:

```

import struct
import sys

def rotate(x, n):
    x &= 0xffffffff
    return ((x << n) | (x >> (32 - n))) & 0xffffffff

def step(s, i, j, k, r):
    s[i] ^= rotate(s[j] + s[k], r)

def quarterround(s, i0, i1, i2, i3):
    step(s, i1, i0, i3, 7)
    step(s, i2, i1, i0, 9)
    step(s, i3, i2, i1, 13)
    step(s, i0, i3, i2, 18)

def rowround(s):
    quarterround(s, 0, 1, 2, 3)
    quarterround(s, 5, 6, 7, 4)
    quarterround(s, 10, 11, 8, 9)
    quarterround(s, 15, 12, 13, 14)

def columnround(s):
    quarterround(s, 0, 4, 8, 12)
    quarterround(s, 5, 9, 13, 1)
    quarterround(s, 10, 14, 2, 6)
    quarterround(s, 15, 3, 7, 11)

def doubleround(s):

```



```

    columnround(s)
    rowround(s)

def rounds(s, n):
    s1 = list(s)
    while n >= 2:
        doubleround(s1)
        n -= 2
    for i in range(16): s[i] = (s[i] + s1[i]) & 0xffffffff

o = struct.unpack('<4I', 'expand 32-byte k')

def block(i, n, k):
    i = i / 64
    i = (i & 0xffffffff, i >> 32)
    s = [0] * 16
    s[:5] = o
    s[1:5] = k[:4]
    s[6:10] = n + i
    s[11:15] = k[4:]
    rounds(s, 20)
    return struct.pack('<16I', *s)

def print_salsa20(l, n, k):
    n = struct.unpack('<2I', n)
    k = struct.unpack('<8I', k)
    for i in xrange(0, l, 64):
        sys.stdout.write(block(i, n, k)[:l-i])

k=[0xdc,0x90,0x8d,0xda,0x0b,0x93,0x44,0xa9
   ,0x53,0x62,0x9b,0x73,0x38,0x20,0x77,0x88
   ,0x80,0xf3,0xce,0xb4,0x21,0xbb,0x61,0xb9
   ,0x1c,0xbd,0x4c,0x3e,0x66,0x25,0x6c,0xe4]
k = ''.join([chr(k[i]) for i in range(32)])

n=[0x82,0x19,0xe0,0x03,0x6b,0x7a,0x0b,0x37]
n = ''.join([chr(n[i]) for i in range(8)])

print_salsa20(4194304,n,k)

```

The output is the same as the 4194304-byte output from the C NaCl program shown earlier.

Testing: Salsa20 specification vs. core_salsa20. The following program uses C NaCl to compute the first Salsa20 example in [7, Section 9]:

```
#include <stdio.h>
```

```

#include "crypto_core_salsa20.h"

unsigned char k[32] = {
    1, 2, 3, 4, 5, 6, 7, 8
, 9, 10, 11, 12, 13, 14, 15, 16
, 201, 202, 203, 204, 205, 206, 207, 208
, 209, 210, 211, 212, 213, 214, 215, 216
} ;

unsigned char in[16] = {
    101, 102, 103, 104, 105, 106, 107, 108
, 109, 110, 111, 112, 113, 114, 115, 116
} ;

unsigned char c[16] = {
    101, 120, 112, 97, 110, 100, 32, 51
, 50, 45, 98, 121, 116, 101, 32, 107
} ;

unsigned char out[64];

main()
{
    int i;
    crypto_core_salsa20(out, in, k, c);
    for (i = 0; i < 64; ++i) {
        if (i > 0) printf(","); else printf(" ");
        printf("%3d", (unsigned int) out[i]);
        if (i % 8 == 7) printf("\n");
    }
    return 0;
}

```

The output of the program is

```

    69, 37, 68, 39, 41, 15, 107, 193
, 255, 139, 122, 6, 170, 233, 217, 98
, 89, 144, 182, 106, 21, 51, 200, 65
, 239, 49, 222, 34, 215, 114, 40, 126
, 104, 197, 7, 225, 197, 153, 31, 2
, 102, 78, 76, 176, 84, 245, 246, 184
, 177, 160, 133, 130, 6, 72, 149, 119
, 192, 195, 132, 236, 234, 103, 246, 74

```

matching the output shown in [7, Section 9].

Testing: core_salsa20 vs. core_hsalsa20. The following program uses C NaCl to compute HSalsa20 on a sample input:

```

#include <stdio.h>
#include "crypto_core_hsalsa20.h"

unsigned char k[32] = {
    0xee,0x30,0x4f,0xca,0x27,0x00,0x8d,0x8c
    ,0x12,0x6f,0x90,0x02,0x79,0x01,0xd8,0x0f
    ,0x7f,0x1d,0x8b,0x8d,0xc9,0x36,0xcf,0x3b
    ,0x9f,0x81,0x96,0x92,0x82,0x7e,0x57,0x77
} ;

unsigned char in[16] = {
    0x81,0x91,0x8e,0xf2,0xa5,0xe0,0xda,0x9b
    ,0x3e,0x90,0x60,0x52,0x1e,0x4b,0xb3,0x52
} ;

unsigned char c[16] = {
    101,120,112, 97,110,100, 32, 51
    , 50, 45, 98,121,116,101, 32,107
} ;

unsigned char out[32];

main()
{
    int i;
    crypto_core_hsalsa20(out,in,k,c);
    for (i = 0;i < 32;++i) {
        printf(",0x%02x", (unsigned int) out[i]);
        if (i % 8 == 7) printf("\n");
    }
    return 0;
}

```

Here is the output of the program:

```

,0xbc,0x1b,0x30,0xfc,0x07,0x2c,0xc1,0x40
,0x75,0xe4,0xba,0xa7,0x31,0xb5,0xa8,0x45
,0xea,0x9b,0x11,0xe9,0xa5,0x19,0x1f,0x94
,0xe1,0x8c,0xba,0x8f,0xd8,0x21,0xa7,0xcd

```

The following program uses C NaCl to compute Salsa20 on the same sample input, and then converts the Salsa20 output to HSalsa20 output:

```

#include <stdio.h>
#include "crypto_core_salsa20.h"

unsigned char k[32] = {

```

```

    0xee,0x30,0x4f,0xca,0x27,0x00,0x8d,0x8c
    ,0x12,0x6f,0x90,0x02,0x79,0x01,0xd8,0x0f
    ,0x7f,0x1d,0x8b,0x8d,0xc9,0x36,0xcf,0x3b
    ,0x9f,0x81,0x96,0x92,0x82,0x7e,0x57,0x77
} ;

unsigned char in[16] = {
    0x81,0x91,0x8e,0xf2,0xa5,0xe0,0xda,0x9b
    ,0x3e,0x90,0x60,0x52,0x1e,0x4b,0xb3,0x52
} ;

unsigned char c[16] = {
    101,120,112, 97,110,100, 32, 51
    , 50, 45, 98,121,116,101, 32,107
} ;

unsigned char out[64];

void print(unsigned char *x,unsigned char *y)
{
    int i;
    unsigned int borrow = 0;
    for (i = 0;i < 4;++i) {
        unsigned int xi = x[i];
        unsigned int yi = y[i];
        printf(",0x%02x",255 & (xi - yi - borrow));
        borrow = (xi < yi + borrow);
    }
}

main()
{
    crypto_core_salsa20(out,in,k,c);
    print(out,c);
    print(out + 20,c + 4); printf("\n");
    print(out + 40,c + 8);
    print(out + 60,c + 12); printf("\n");
    print(out + 24,in);
    print(out + 28,in + 4); printf("\n");
    print(out + 32,in + 8);
    print(out + 36,in + 12); printf("\n");
    return 0;
}

```

This program produces the same output as the previous program.

9 Plaintext, ciphertext, and authenticator

To encrypt a packet $m \in \{0, 1, \dots, 255\}^{\{0, 1, \dots, 2^{70} - 32\}}$ using the packet's nonce $n \in \{0, 1, \dots, 255\}^{24}$ and the shared secret $k \in \{0, 1, \dots, 255\}^{32}$, Alice xors the packet with *part of* the long stream computed in the previous section. Alice then uses a different part of the long stream to authenticate the ciphertext. Alice's boxed packet is the authenticator followed by the ciphertext.

Specifically, write the nonce n as (n_1, n_2) with $n_1 \in \{0, 1, \dots, 255\}^{16}$ and $n_2 \in \{0, 1, \dots, 255\}^8$, and write $\text{Salsa20}(\text{HSalsa20}(\text{HSalsa20}(k, 0), n_1), n_2)$ as (r, s, t, \dots) where $r, s \in \{0, 1, \dots, 255\}^{16}$ and $\text{len } t = \text{len } m$. Define $c = m \oplus t \in \{0, 1, \dots, 255\}^{\text{len } m}$ and $a = \text{Poly1305}(\text{ClampP}(r), c, s) \in \{0, 1, \dots, 255\}^{16}$. The boxed packet is then $(a, c) \in \{0, 1, \dots, 255\}^{16 + \text{len } m}$.

This section defines Poly1305 and ClampP. Some of the definitions here are copied from the original Poly1305 specification [6]; $\text{Poly1305}(r, c, s)$ here is $\text{Poly1305}_r(c, s)$ in the notation of [6].

The ClampP function. The function

$$\text{ClampP} : \{0, 1, \dots, 255\}^{16} \rightarrow \{0, 1, \dots, 255\}^{16}$$

maps $(r_0, r_1, \dots, r_{15})$ to

$$\begin{aligned} & (r_0, & r_1, & r_2, & r_3 \bmod 16, \\ & r_4 - (r_4 \bmod 4), & r_5, & r_6, & r_7 \bmod 16, \\ & r_8 - (r_8 \bmod 4), & r_9, & r_{10}, & r_{11} \bmod 16, \\ & r_{12} - (r_{12} \bmod 4), & r_{13}, & r_{14}, & r_{15} \bmod 16). \end{aligned}$$

The Poly1305 function. Fix $\ell \in \{0, 1, \dots, 2^{70} - 32\}$, fix $c \in \{0, 1, \dots, 255\}^\ell$, fix $R \in \{0, 1, \dots, 2^{128} - 1\}$, and fix $S \in \{0, 1, \dots, 2^{128} - 1\}$. Write $q = \lceil \ell / 16 \rceil$. Write c as $(c[0], c[1], \dots, c[\ell - 1])$. Define $C_1, C_2, \dots, C_q \in \{1, 2, 3, \dots, 2^{129}\}$ as follows: if $1 \leq i \leq \lfloor \ell / 16 \rfloor$ then

$$C_i = c[16i - 16] + 2^8 c[16i - 15] + 2^{16} c[16i - 14] + \dots + 2^{120} c[16i - 1] + 2^{128};$$

if ℓ is not a multiple of 16 then

$$C_q = c[16q - 16] + 2^8 c[16q - 15] + \dots + 2^{8(\ell \bmod 16) - 8} c[\ell - 1] + 2^{8(\ell \bmod 16)}.$$

In other words: Pad each 16-byte chunk of the ciphertext to 17 bytes by appending a 1. If the ciphertext 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.

Now $\text{Poly1305}(R, c, S) = A$ where

$$A = (((C_1 R^q + C_2 R^{q-1} + \dots + C_q R^1) \bmod 2^{130} - 5) + S) \bmod 2^{128}.$$

Here \underline{A} means the 16-byte string $(A \bmod 256, \lfloor A/256 \rfloor \bmod 256, \dots)$; \underline{R} and \underline{S} are defined in the same way.

Security notes. The constructions in this section—xor for encryption and Poly1305 for authentication—are *provably* secure. If the attacker cannot distinguish the stream (r, s, t) from a uniform random string then the attacker learns nothing about the original packet m , aside from its length, and has negligible chance of replacing the boxed packet (a, c) with a different packet (a', c') that satisfies $a' = \text{Poly1305}(\text{ClampP}(r), c', s)$. Of course, this guarantee says nothing about an attacker who *can* distinguish (r, s, t) from a uniform random string—for example, an attacker who uses a quantum computer to break elliptic-curve cryptography has as much power as Alice and Bob.

A security proof for Poly1305 appears in [6]. The proof shows that if packets are limited to L bytes then the attacker’s success chance for a forgery attempt (a', c') is at most $8\lceil L/16 \rceil / 2^{106}$. Here are some of the critical points in the proof: $2^{130} - 5$ is prime; $\text{ClampP}(r)$ is uniformly distributed among 2^{106} possibilities; and distinct strings c produce distinct polynomials $C_1x^q + C_2x^{q-1} + \dots + C_qx^1$ modulo $2^{130} - 5$.

What happens if an attacker is astonishingly lucky and succeeds at a forgery attempt? Presumably this success will be visible from the receiver’s behavior. The attacker can then, by polynomial root-finding, easily determine $\text{ClampP}(r)$ and s , or at worst a short list of possibilities for $\text{ClampP}(r)$ and s , allowing the attacker to generate “re-forgeries” (a'', c'') under the same nonce. However, if the receiver follows the standard practice of insisting on a strictly increasing sequence of nonces, then the receiver will reject all of these “re-forgeries,” as pointed out in 2005 by Nyberg, Gilbert, and Robshaw and independently in 2006 by Lange. See [19] and [9, Section 2.5].

If r were reused from one nonce to another, with s generated anew for each nonce, then the first forgery would still be difficult (as pointed out by Wegman and Carter in [24, Section 4]), but after seeing a successful forgery the attacker would be able to generate “re-forgeries” under *other* nonces. If >100 -bit security were scaled down to much lower security then the attacker could reasonably hope for this situation to occur. Many authentication systems in the literature have this problem. The following comment appears in [5, Section 8] and was already online in 2000:

Some writers claim that forgery probabilities around $1/2^{32}$ are adequate for most applications. The attacker’s cost of 2^{32} forgery attempts, they say, is much larger than the attacker’s benefit from forging a single message. Unfortunately, even if all attackers acted on the basis of rational economic analyses, this argument would be wrong, because it wildly underestimates the attacker’s benefit. *In a typical authentication system, as soon as the attacker is lucky enough to succeed at a few forgeries, he can immediately figure out enough secret information to let him forge messages of his choice.* (This does not contradict the information-theoretic security expressed by Theorem 8.2; the attacker is gaining information

from the receiver, not from the sender.) It is crucial for the forgery probability to be so small that attackers have no hope.

(Emphasis added.) Detailed explanations of various re-forgery attacks appeared in [16], [15], and [11].

Attacks of that type do not apply to Poly1305 as used in NaCl. There is a new Poly1305 key (r, s) for each nonce; the standard security conjecture for Salsa20 implies that the keys (r, s) for different nonces are indistinguishable from independent uniform random keys. More importantly, the >100 -bit security level of Poly1305 prevents forgery attempts from succeeding in the first place.

10 Example of the plaintext, ciphertext, and authenticator

The following program starts from Section 3's example of Alice's secret key a , Section 4's example of Bob's public key B , Section 8's example of a nonce n , and a sample 131-byte packet, and uses C NaCl to compute the corresponding boxed packet:

```
#include <stdio.h>
#include "crypto_box_curve25519xsalsa20poly1305.h"

unsigned char alicesk[32] = {
    0x77,0x07,0x6d,0x0a,0x73,0x18,0xa5,0x7d
    ,0x3c,0x16,0xc1,0x72,0x51,0xb2,0x66,0x45
    ,0xdf,0x4c,0x2f,0x87,0xeb,0xc0,0x99,0x2a
    ,0xb1,0x77,0xfb,0xa5,0x1d,0xb9,0x2c,0x2a
} ;

unsigned char bobpk[32] = {
    0xde,0x9e,0xdb,0x7d,0x7b,0x7d,0xc1,0xb4
    ,0xd3,0x5b,0x61,0xc2,0xec,0xe4,0x35,0x37
    ,0x3f,0x83,0x43,0xc8,0x5b,0x78,0x67,0x4d
    ,0xad,0xfc,0x7e,0x14,0x6f,0x88,0x2b,0x4f
} ;

unsigned char nonce[24] = {
    0x69,0x69,0x6e,0xe9,0x55,0xb6,0x2b,0x73
    ,0xcd,0x62,0xbd,0xa8,0x75,0xfc,0x73,0xd6
    ,0x82,0x19,0xe0,0x03,0x6b,0x7a,0x0b,0x37
} ;

// API requires first 32 bytes to be 0
unsigned char m[163] = {
    0, 0, 0, 0, 0, 0, 0, 0
    , 0, 0, 0, 0, 0, 0, 0, 0
```

```

, 0, 0, 0, 0, 0, 0, 0, 0
, 0, 0, 0, 0, 0, 0, 0, 0
, 0xbe, 0x07, 0x5f, 0xc5, 0x3c, 0x81, 0xf2, 0xd5
, 0xcf, 0x14, 0x13, 0x16, 0xeb, 0xeb, 0x0c, 0x7b
, 0x52, 0x28, 0xc5, 0x2a, 0x4c, 0x62, 0xcb, 0xd4
, 0x4b, 0x66, 0x84, 0x9b, 0x64, 0x24, 0x4f, 0xfc
, 0xe5, 0xec, 0xba, 0xaf, 0x33, 0xbd, 0x75, 0x1a
, 0x1a, 0xc7, 0x28, 0xd4, 0x5e, 0x6c, 0x61, 0x29
, 0x6c, 0xdc, 0x3c, 0x01, 0x23, 0x35, 0x61, 0xf4
, 0x1d, 0xb6, 0x6c, 0xce, 0x31, 0x4a, 0xdb, 0x31
, 0x0e, 0x3b, 0xe8, 0x25, 0x0c, 0x46, 0xf0, 0x6d
, 0xce, 0xea, 0x3a, 0x7f, 0xa1, 0x34, 0x80, 0x57
, 0xe2, 0xf6, 0x55, 0x6a, 0xd6, 0xb1, 0x31, 0x8a
, 0x02, 0x4a, 0x83, 0x8f, 0x21, 0xaf, 0x1f, 0xde
, 0x04, 0x89, 0x77, 0xeb, 0x48, 0xf5, 0x9f, 0xfd
, 0x49, 0x24, 0xca, 0x1c, 0x60, 0x90, 0x2e, 0x52
, 0xf0, 0xa0, 0x89, 0xbc, 0x76, 0x89, 0x70, 0x40
, 0xe0, 0x82, 0xf9, 0x37, 0x76, 0x38, 0x48, 0x64
, 0x5e, 0x07, 0x05
} ;

```

```

unsigned char c[163];

```

```

main()
{
    int i;
    crypto_box_curve25519xsalsa20poly1305(
        c, m, 163, nonce, bobpk, alicesk
    );
    for (i = 16; i < 163; ++i) {
        printf(", 0x%02x", (unsigned int) c[i]);
        if (i % 8 == 7) printf("\n");
    }
    printf("\n");
    return 0;
}

```

The program prints a 147-byte boxed packet:

```

, 0xf3, 0xff, 0xc7, 0x70, 0x3f, 0x94, 0x00, 0xe5
, 0x2a, 0x7d, 0xfb, 0x4b, 0x3d, 0x33, 0x05, 0xd9
, 0x8e, 0x99, 0x3b, 0x9f, 0x48, 0x68, 0x12, 0x73
, 0xc2, 0x96, 0x50, 0xba, 0x32, 0xfc, 0x76, 0xce
, 0x48, 0x33, 0x2e, 0xa7, 0x16, 0x4d, 0x96, 0xa4
, 0x47, 0x6f, 0xb8, 0xc5, 0x31, 0xa1, 0x18, 0x6a
, 0xc0, 0xdf, 0xc1, 0x7c, 0x98, 0xdc, 0xe8, 0x7b
, 0x4d, 0xa7, 0xf0, 0x11, 0xec, 0x48, 0xc9, 0x72

```



```
,0x71,0xd2,0xc2,0x0f,0x9b,0x92,0x8f,0xe2
,0x27,0x0d,0x6f,0xb8,0x63,0xd5,0x17,0x38
,0xb4,0x8e,0xee,0xe3,0x14,0xa7,0xcc,0x8a
,0xb9,0x32,0x16,0x45,0x48,0xe5,0x26,0xae
,0x90,0x22,0x43,0x68,0x51,0x7a,0xcf,0xea
,0xbd,0x6b,0xb3,0x73,0x2b,0xc0,0xe9,0xda
,0x99,0x83,0x2b,0x61,0xca,0x01,0xb6,0xde
,0x56,0x24,0x4a,0x9e,0x88,0xd5,0xf9,0xb3
,0x79,0x73,0xf6,0x22,0xa4,0x3d,0x14,0xa6
,0x59,0x9b,0x1f,0x65,0x4c,0xb4,0x5a,0x74
,0xe3,0x55,0xa5
```

The following program starts from Section 4's example of Bob's secret key *b*, Section 3's example of Alice's public key *A*, Section 8's example of a nonce *n*, and the 147-byte boxed packet shown above, and uses C NaCl to open the box:

```
#include <stdio.h>
#include "crypto_box_curve25519xsalsa20poly1305.h"

unsigned char bobsk[32] = {
    0x5d,0xab,0x08,0x7e,0x62,0x4a,0x8a,0x4b
,0x79,0xe1,0x7f,0x8b,0x83,0x80,0x0e,0xe6
,0x6f,0x3b,0xb1,0x29,0x26,0x18,0xb6,0xfd
,0x1c,0x2f,0x8b,0x27,0xff,0x88,0xe0,0xeb
} ;

unsigned char alicepk[32] = {
    0x85,0x20,0xf0,0x09,0x89,0x30,0xa7,0x54
,0x74,0x8b,0x7d,0xdc,0xb4,0x3e,0xf7,0x5a
,0x0d,0xbf,0x3a,0x0d,0x26,0x38,0x1a,0xf4
,0xeb,0xa4,0xa9,0x8e,0xaa,0x9b,0x4e,0x6a
} ;

unsigned char nonce[24] = {
    0x69,0x69,0x6e,0xe9,0x55,0xb6,0x2b,0x73
,0xcd,0x62,0xbd,0xa8,0x75,0xfc,0x73,0xd6
,0x82,0x19,0xe0,0x03,0x6b,0x7a,0x0b,0x37
} ;

// API requires first 16 bytes to be 0
unsigned char c[163] = {
    0, 0, 0, 0, 0, 0, 0, 0
, 0, 0, 0, 0, 0, 0, 0, 0
,0xf3,0xff,0xc7,0x70,0x3f,0x94,0x00,0xe5
,0x2a,0x7d,0xfb,0x4b,0x3d,0x33,0x05,0xd9
,0x8e,0x99,0x3b,0x9f,0x48,0x68,0x12,0x73
,0xc2,0x96,0x50,0xba,0x32,0xfc,0x76,0xce
```

```

,0x48,0x33,0x2e,0xa7,0x16,0x4d,0x96,0xa4
,0x47,0x6f,0xb8,0xc5,0x31,0xa1,0x18,0x6a
,0xc0,0xdf,0xc1,0x7c,0x98,0xdc,0xe8,0x7b
,0x4d,0xa7,0xf0,0x11,0xec,0x48,0xc9,0x72
,0x71,0xd2,0xc2,0x0f,0x9b,0x92,0x8f,0xe2
,0x27,0x0d,0x6f,0xb8,0x63,0xd5,0x17,0x38
,0xb4,0x8e,0xee,0xe3,0x14,0xa7,0xcc,0x8a
,0xb9,0x32,0x16,0x45,0x48,0xe5,0x26,0xae
,0x90,0x22,0x43,0x68,0x51,0x7a,0xcf,0xea
,0xbd,0x6b,0xb3,0x73,0x2b,0xc0,0xe9,0xda
,0x99,0x83,0x2b,0x61,0xca,0x01,0xb6,0xde
,0x56,0x24,0x4a,0x9e,0x88,0xd5,0xf9,0xb3
,0x79,0x73,0xf6,0x22,0xa4,0x3d,0x14,0xa6
,0x59,0x9b,0x1f,0x65,0x4c,0xb4,0x5a,0x74
,0xe3,0x55,0xa5
} ;

```

```

unsigned char m[163];

```

```

main()
{
    int i;
    if (crypto_box_curve25519xsalsa20poly1305_open(
        m,c,163,nonce,alicepk,bobsk
    ) == 0) {
        for (i = 32;i < 163;++i) {
            printf("0x%02x", (unsigned int) m[i]);
            if (i % 8 == 7) printf("\n");
        }
        printf("\n");
    }
    return 0;
}

```

The program prints the original 131-byte packet:

```

,0xbe,0x07,0x5f,0xc5,0x3c,0x81,0xf2,0xd5
,0xcf,0x14,0x13,0x16,0xeb,0xeb,0x0c,0x7b
,0x52,0x28,0xc5,0x2a,0x4c,0x62,0xcb,0xd4
,0x4b,0x66,0x84,0x9b,0x64,0x24,0x4f,0xfc
,0xe5,0xec,0xba,0xaf,0x33,0xbd,0x75,0x1a
,0x1a,0xc7,0x28,0xd4,0x5e,0x6c,0x61,0x29
,0x6c,0xdc,0x3c,0x01,0x23,0x35,0x61,0xf4
,0x1d,0xb6,0x6c,0xce,0x31,0x4a,0xdb,0x31
,0x0e,0x3b,0xe8,0x25,0x0c,0x46,0xf0,0x6d
,0xce,0xea,0x3a,0x7f,0xa1,0x34,0x80,0x57
,0xe2,0xf6,0x55,0x6a,0xd6,0xb1,0x31,0x8a

```

```
,0x02,0x4a,0x83,0x8f,0x21,0xaf,0x1f,0xde
,0x04,0x89,0x77,0xeb,0x48,0xf5,0x9f,0xfd
,0x49,0x24,0xca,0x1c,0x60,0x90,0x2e,0x52
,0xf0,0xa0,0x89,0xbc,0x76,0x89,0x70,0x40
,0xe0,0x82,0xf9,0x37,0x76,0x38,0x48,0x64
,0x5e,0x07,0x05
```

Testing: box vs. secretbox. The following program computes the same 147-byte boxed packet, but starts from the first-level key k_1 computed in Section 8:

```
#include <stdio.h>
#include "crypto_secretbox_xsalsa20poly1305.h"

unsigned char firstkey[32] = {
    0x1b,0x27,0x55,0x64,0x73,0xe9,0x85,0xd4
, 0x62,0xcd,0x51,0x19,0x7a,0x9a,0x46,0xc7
, 0x60,0x09,0x54,0x9e,0xac,0x64,0x74,0xf2
, 0x06,0xc4,0xee,0x08,0x44,0xf6,0x83,0x89
} ;

unsigned char nonce[24] = {
    0x69,0x69,0x6e,0xe9,0x55,0xb6,0x2b,0x73
, 0xcd,0x62,0xbd,0xa8,0x75,0xfc,0x73,0xd6
, 0x82,0x19,0xe0,0x03,0x6b,0x7a,0x0b,0x37
} ;

// API requires first 32 bytes to be 0
unsigned char m[163] = {
    0, 0, 0, 0, 0, 0, 0, 0
, 0, 0, 0, 0, 0, 0, 0, 0
, 0, 0, 0, 0, 0, 0, 0, 0
, 0, 0, 0, 0, 0, 0, 0, 0
,0xbe,0x07,0x5f,0xc5,0x3c,0x81,0xf2,0xd5
,0xcf,0x14,0x13,0x16,0xeb,0xeb,0x0c,0x7b
,0x52,0x28,0xc5,0x2a,0x4c,0x62,0xcb,0xd4
,0x4b,0x66,0x84,0x9b,0x64,0x24,0x4f,0xfc
,0xe5,0xec,0xba,0xaf,0x33,0xbd,0x75,0x1a
,0x1a,0xc7,0x28,0xd4,0x5e,0x6c,0x61,0x29
,0x6c,0xdc,0x3c,0x01,0x23,0x35,0x61,0xf4
,0x1d,0xb6,0x6c,0xce,0x31,0x4a,0xdb,0x31
,0x0e,0x3b,0xe8,0x25,0x0c,0x46,0xf0,0x6d
,0xce,0xea,0x3a,0x7f,0xa1,0x34,0x80,0x57
,0xe2,0xf6,0x55,0x6a,0xd6,0xb1,0x31,0x8a
,0x02,0x4a,0x83,0x8f,0x21,0xaf,0x1f,0xde
,0x04,0x89,0x77,0xeb,0x48,0xf5,0x9f,0xfd
```

```

,0x49,0x24,0xca,0x1c,0x60,0x90,0x2e,0x52
,0xf0,0xa0,0x89,0xbc,0x76,0x89,0x70,0x40
,0xe0,0x82,0xf9,0x37,0x76,0x38,0x48,0x64
,0x5e,0x07,0x05
} ;

unsigned char c[163];

main()
{
    int i;
    crypto_secretbox_xsalsa20poly1305(
        c,m,163,nonce,firstkey
    );
    for (i = 16;i < 163;++i) {
        printf(",0x%02x",(unsigned int) c[i]);
        if (i % 8 == 7) printf("\n");
    }
    printf("\n");
    return 0;
}

```

The following program opens the same box, again starting from the first-level key k_1 :

```

#include <stdio.h>
#include "crypto_secretbox_xsalsa20poly1305.h"

unsigned char firstkey[32] = {
    0x1b,0x27,0x55,0x64,0x73,0xe9,0x85,0xd4
,0x62,0xcd,0x51,0x19,0x7a,0x9a,0x46,0xc7
,0x60,0x09,0x54,0x9e,0xac,0x64,0x74,0xf2
,0x06,0xc4,0xee,0x08,0x44,0xf6,0x83,0x89
} ;

unsigned char nonce[24] = {
    0x69,0x69,0x6e,0xe9,0x55,0xb6,0x2b,0x73
,0xcd,0x62,0xbd,0xa8,0x75,0xfc,0x73,0xd6
,0x82,0x19,0xe0,0x03,0x6b,0x7a,0x0b,0x37
} ;

// API requires first 16 bytes to be 0
unsigned char c[163] = {
    0, 0, 0, 0, 0, 0, 0, 0
, 0, 0, 0, 0, 0, 0, 0, 0
,0xf3,0xff,0xc7,0x70,0x3f,0x94,0x00,0xe5
,0x2a,0x7d,0xfb,0x4b,0x3d,0x33,0x05,0xd9

```

```

,0x8e,0x99,0x3b,0x9f,0x48,0x68,0x12,0x73
,0xc2,0x96,0x50,0xba,0x32,0xfc,0x76,0xce
,0x48,0x33,0x2e,0xa7,0x16,0x4d,0x96,0xa4
,0x47,0x6f,0xb8,0xc5,0x31,0xa1,0x18,0x6a
,0xc0,0xdf,0xc1,0x7c,0x98,0xdc,0xe8,0x7b
,0x4d,0xa7,0xf0,0x11,0xec,0x48,0xc9,0x72
,0x71,0xd2,0xc2,0x0f,0x9b,0x92,0x8f,0xe2
,0x27,0x0d,0x6f,0xb8,0x63,0xd5,0x17,0x38
,0xb4,0x8e,0xee,0xe3,0x14,0xa7,0xcc,0x8a
,0xb9,0x32,0x16,0x45,0x48,0xe5,0x26,0xae
,0x90,0x22,0x43,0x68,0x51,0x7a,0xcf,0xea
,0xbd,0x6b,0xb3,0x73,0x2b,0xc0,0xe9,0xda
,0x99,0x83,0x2b,0x61,0xca,0x01,0xb6,0xde
,0x56,0x24,0x4a,0x9e,0x88,0xd5,0xf9,0xb3
,0x79,0x73,0xf6,0x22,0xa4,0x3d,0x14,0xa6
,0x59,0x9b,0x1f,0x65,0x4c,0xb4,0x5a,0x74
,0xe3,0x55,0xa5
} ;

unsigned char m[163];

main()
{
    int i;
    if (crypto_secretbox_xsalsa20poly1305_open(
        m,c,163,nonce,firstkey
    ) == 0) {
        for (i = 32;i < 163;++i) {
            printf(",0x%02x",(unsigned int) m[i]);
            if (i % 8 == 7) printf("\n");
        }
        printf("\n");
    }
    return 0;
}

```

Testing: secretbox vs. stream. The following program starts from the first-level key k_1 shown above, computes the first 163 bytes of the corresponding stream as in Section 8, skips the first 32 bytes, and xors the remaining bytes with the 131-byte packet shown above:

```

#include <stdio.h>
#include "crypto_stream_xsalsa20.h"

unsigned char firstkey[32] = {
    0x1b,0x27,0x55,0x64,0x73,0xe9,0x85,0xd4

```

```
,0x62,0xcd,0x51,0x19,0x7a,0x9a,0x46,0xc7
,0x60,0x09,0x54,0x9e,0xac,0x64,0x74,0xf2
,0x06,0xc4,0xee,0x08,0x44,0xf6,0x83,0x89
} ;
```

```
unsigned char nonce[24] = {
    0x69,0x69,0x6e,0xe9,0x55,0xb6,0x2b,0x73
,0xcd,0x62,0xbd,0xa8,0x75,0xfc,0x73,0xd6
,0x82,0x19,0xe0,0x03,0x6b,0x7a,0x0b,0x37
} ;
```

```
unsigned char m[163] = {
    0, 0, 0, 0, 0, 0, 0, 0
, 0, 0, 0, 0, 0, 0, 0, 0
, 0, 0, 0, 0, 0, 0, 0, 0
, 0, 0, 0, 0, 0, 0, 0, 0
,0xbe,0x07,0x5f,0xc5,0x3c,0x81,0xf2,0xd5
,0xcf,0x14,0x13,0x16,0xeb,0xeb,0x0c,0x7b
,0x52,0x28,0xc5,0x2a,0x4c,0x62,0xcb,0xd4
,0x4b,0x66,0x84,0x9b,0x64,0x24,0x4f,0xfc
,0xe5,0xec,0xba,0xaf,0x33,0xbd,0x75,0x1a
,0x1a,0xc7,0x28,0xd4,0x5e,0x6c,0x61,0x29
,0x6c,0xdc,0x3c,0x01,0x23,0x35,0x61,0xf4
,0x1d,0xb6,0x6c,0xce,0x31,0x4a,0xdb,0x31
,0x0e,0x3b,0xe8,0x25,0x0c,0x46,0xf0,0x6d
,0xce,0xea,0x3a,0x7f,0xa1,0x34,0x80,0x57
,0xe2,0xf6,0x55,0x6a,0xd6,0xb1,0x31,0x8a
,0x02,0x4a,0x83,0x8f,0x21,0xaf,0x1f,0xde
,0x04,0x89,0x77,0xeb,0x48,0xf5,0x9f,0xfd
,0x49,0x24,0xca,0x1c,0x60,0x90,0x2e,0x52
,0xf0,0xa0,0x89,0xbc,0x76,0x89,0x70,0x40
,0xe0,0x82,0xf9,0x37,0x76,0x38,0x48,0x64
,0x5e,0x07,0x05
} ;
```

```
unsigned char c[163];
```

```
main()
{
    int i;
    crypto_stream_xsalsa20_xor(c,m,163,nonce,firstkey);
    for (i = 32;i < 163;++i) {
        printf("0x%02x",(unsigned int) c[i]);
        if (i % 8 == 7) printf("\n");
    }
}
```

```

    printf("\n");
    return 0;
}

```

This program prints

```

,0x8e,0x99,0x3b,0x9f,0x48,0x68,0x12,0x73
,0xc2,0x96,0x50,0xba,0x32,0xfc,0x76,0xce
,0x48,0x33,0x2e,0xa7,0x16,0x4d,0x96,0xa4
,0x47,0x6f,0xb8,0xc5,0x31,0xa1,0x18,0x6a
,0xc0,0xdf,0xc1,0x7c,0x98,0xdc,0xe8,0x7b
,0x4d,0xa7,0xf0,0x11,0xec,0x48,0xc9,0x72
,0x71,0xd2,0xc2,0x0f,0x9b,0x92,0x8f,0xe2
,0x27,0x0d,0x6f,0xb8,0x63,0xd5,0x17,0x38
,0xb4,0x8e,0xee,0xe3,0x14,0xa7,0xcc,0x8a
,0xb9,0x32,0x16,0x45,0x48,0xe5,0x26,0xae
,0x90,0x22,0x43,0x68,0x51,0x7a,0xcf,0xea
,0xbd,0x6b,0xb3,0x73,0x2b,0xc0,0xe9,0xda
,0x99,0x83,0x2b,0x61,0xca,0x01,0xb6,0xde
,0x56,0x24,0x4a,0x9e,0x88,0xd5,0xf9,0xb3
,0x79,0x73,0xf6,0x22,0xa4,0x3d,0x14,0xa6
,0x59,0x9b,0x1f,0x65,0x4c,0xb4,0x5a,0x74
,0xe3,0x55,0xa5

```

matching the final 131 bytes of the 147-byte boxed packet shown above.

Testing: secretbox vs. onetimeauth. The following program starts from the first-level key k_1 shown above and prints the first 32 bytes of the corresponding stream:

```

#include <stdio.h>
#include "crypto_stream_xsalsa20.h"

unsigned char firstkey[32] = {
    0x1b,0x27,0x55,0x64,0x73,0xe9,0x85,0xd4
,0x62,0xcd,0x51,0x19,0x7a,0x9a,0x46,0xc7
,0x60,0x09,0x54,0x9e,0xac,0x64,0x74,0xf2
,0x06,0xc4,0xee,0x08,0x44,0xf6,0x83,0x89
} ;

unsigned char nonce[24] = {
    0x69,0x69,0x6e,0xe9,0x55,0xb6,0x2b,0x73
,0xcd,0x62,0xbd,0xa8,0x75,0xfc,0x73,0xd6
,0x82,0x19,0xe0,0x03,0x6b,0x7a,0x0b,0x37
} ;

unsigned char rs[32];

```

```

main()
{
    int i;
    crypto_stream_xsalsa20(rs,32,nonce,firstkey);
    for (i = 0;i < 32;++i) {
        printf(",0x%02x",(unsigned int) rs[i]);
        if (i % 8 == 7) printf("\n");
    }
    return 0;
}

```

The output of the program is a Poly1305 key (r, s) :

```

,0xee,0xa6,0xa7,0x25,0x1c,0x1e,0x72,0x91
,0x6d,0x11,0xc2,0xcb,0x21,0x4d,0x3c,0x25
,0x25,0x39,0x12,0x1d,0x8e,0x23,0x4e,0x65
,0x2d,0x65,0x1f,0xa4,0xc8,0xcf,0xf8,0x80

```

The following program starts from this Poly1305 key (r, s) and the 131-byte suffix c of the boxed packet shown above, and uses C NaCl to compute $\text{Poly1305}(\text{ClampP}(r), c, s)$:

```

#include <stdio.h>
#include "crypto_onetimeauth_poly1305.h"

unsigned char rs[32] = {
    0xee,0xa6,0xa7,0x25,0x1c,0x1e,0x72,0x91
,0x6d,0x11,0xc2,0xcb,0x21,0x4d,0x3c,0x25
,0x25,0x39,0x12,0x1d,0x8e,0x23,0x4e,0x65
,0x2d,0x65,0x1f,0xa4,0xc8,0xcf,0xf8,0x80
} ;

unsigned char c[131] = {
    0x8e,0x99,0x3b,0x9f,0x48,0x68,0x12,0x73
,0xc2,0x96,0x50,0xba,0x32,0xfc,0x76,0xce
,0x48,0x33,0x2e,0xa7,0x16,0x4d,0x96,0xa4
,0x47,0x6f,0xb8,0xc5,0x31,0xa1,0x18,0x6a
,0xc0,0xdf,0xc1,0x7c,0x98,0xdc,0xe8,0x7b
,0x4d,0xa7,0xf0,0x11,0xec,0x48,0xc9,0x72
,0x71,0xd2,0xc2,0x0f,0x9b,0x92,0x8f,0xe2
,0x27,0x0d,0x6f,0xb8,0x63,0xd5,0x17,0x38
,0xb4,0x8e,0xee,0xe3,0x14,0xa7,0xcc,0x8a
,0xb9,0x32,0x16,0x45,0x48,0xe5,0x26,0xae
,0x90,0x22,0x43,0x68,0x51,0x7a,0xcf,0xea
,0xbd,0x6b,0xb3,0x73,0x2b,0xc0,0xe9,0xda
,0x99,0x83,0x2b,0x61,0xca,0x01,0xb6,0xde
,0x56,0x24,0x4a,0x9e,0x88,0xd5,0xf9,0xb3

```



```

,0x79,0x73,0xf6,0x22,0xa4,0x3d,0x14,0xa6
,0x59,0x9b,0x1f,0x65,0x4c,0xb4,0x5a,0x74
,0xe3,0x55,0xa5
} ;

unsigned char a[16];

main()
{
    int i;
    crypto_onetimeauth_poly1305(a,c,131,rs);
    for (i = 0;i < 16;++i) {
        printf(",0x%02x",(unsigned int) a[i]);
        if (i % 8 == 7) printf("\n");
    }
    return 0;
}

```

The program prints

```

,0xf3,0xff,0xc7,0x70,0x3f,0x94,0x00,0xe5
,0x2a,0x7d,0xfb,0x4b,0x3d,0x33,0x05,0xd9

```

matching the first 16 bytes of the boxed packet shown above.

Testing: C++ vs. onetimeauth. The following C++ program starts from the same Poly1305 key (r,s) and the same c as above, and uses GMP (through GMP's C++ interface) to compute $\text{Poly1305}(\text{ClampP}(r),c,s)$:

```

#include <stdio.h>
#include <gmpxx.h>

void poly1305_gmpxx(unsigned char *out,
    const unsigned char *r,
    const unsigned char *s,
    const unsigned char *m,unsigned int l)
{
    unsigned int j;
    mpz_class rbar = 0;
    for (j = 0;j < 16;++j) {
        mpz_class rj = r[j];
        if (j % 4 == 3) rj = r[j] % 16;
        if (j == 4) rj = r[j] & 252;
        if (j == 8) rj = r[j] & 252;
        if (j == 12) rj = r[j] & 252;
        rbar += rj << (8 * j);
    }
    mpz_class h = 0;

```

```

mpz_class p = (((mpz_class) 1) << 130) - 5;
while (l > 0) {
    mpz_class c = 0;
    for (j = 0; (j < 16) && (j < l); ++j)
        c += ((mpz_class) m[j]) << (8 * j);
    c += ((mpz_class) 1) << (8 * j);
    m += j; l -= j;
    h = ((h + c) * rbar) % p;
}
for (j = 0; j < 16; ++j)
    h += ((mpz_class) s[j]) << (8 * j);
for (j = 0; j < 16; ++j) {
    mpz_class c = h % 256;
    h >>= 8;
    out[j] = c.get_ui();
}
}

```

```

unsigned char rs[32] = {
    0xee, 0xa6, 0xa7, 0x25, 0x1c, 0x1e, 0x72, 0x91
, 0x6d, 0x11, 0xc2, 0xcb, 0x21, 0x4d, 0x3c, 0x25
, 0x25, 0x39, 0x12, 0x1d, 0x8e, 0x23, 0x4e, 0x65
, 0x2d, 0x65, 0x1f, 0xa4, 0xc8, 0xcf, 0xf8, 0x80
} ;

```

```

unsigned char c[131] = {
    0x8e, 0x99, 0x3b, 0x9f, 0x48, 0x68, 0x12, 0x73
, 0xc2, 0x96, 0x50, 0xba, 0x32, 0xfc, 0x76, 0xce
, 0x48, 0x33, 0x2e, 0xa7, 0x16, 0x4d, 0x96, 0xa4
, 0x47, 0x6f, 0xb8, 0xc5, 0x31, 0xa1, 0x18, 0x6a
, 0xc0, 0xdf, 0xc1, 0x7c, 0x98, 0xdc, 0xe8, 0x7b
, 0x4d, 0xa7, 0xf0, 0x11, 0xec, 0x48, 0xc9, 0x72
, 0x71, 0xd2, 0xc2, 0x0f, 0x9b, 0x92, 0x8f, 0xe2
, 0x27, 0x0d, 0x6f, 0xb8, 0x63, 0xd5, 0x17, 0x38
, 0xb4, 0x8e, 0xee, 0xe3, 0x14, 0xa7, 0xcc, 0x8a
, 0xb9, 0x32, 0x16, 0x45, 0x48, 0xe5, 0x26, 0xae
, 0x90, 0x22, 0x43, 0x68, 0x51, 0x7a, 0xcf, 0xea
, 0xbd, 0x6b, 0xb3, 0x73, 0x2b, 0xc0, 0xe9, 0xda
, 0x99, 0x83, 0x2b, 0x61, 0xca, 0x01, 0xb6, 0xde
, 0x56, 0x24, 0x4a, 0x9e, 0x88, 0xd5, 0xf9, 0xb3
, 0x79, 0x73, 0xf6, 0x22, 0xa4, 0x3d, 0x14, 0xa6
, 0x59, 0x9b, 0x1f, 0x65, 0x4c, 0xb4, 0x5a, 0x74
, 0xe3, 0x55, 0xa5
} ;

```

```

unsigned char a[16];

main()
{
    int i;
    poly1305_gmpxx(a,rs,rs + 16,c,131);
    for (i = 0;i < 16;++i) {
        printf("0x%02x",(unsigned int) a[i]);
        if (i % 8 == 7) printf("\n");
    }
    return 0;
}

```

The program prints

```

,0xf3,0xff,0xc7,0x70,0x3f,0x94,0x00,0xe5
,0x2a,0x7d,0xfb,0x4b,0x3d,0x33,0x05,0xd9

```

matching the output of the previous program.

References

1. — (no editor), *20th annual symposium on foundations of computer science*, IEEE Computer Society, New York, 1979. MR 82a:68004. See [23].
2. Jean-Philippe Aumasson, Simon Fischer, Shahram Khazaei, Willi Meier, Christian Rechberger, *New features of Latin dances: analysis of Salsa, ChaCha, and Rumba* (2007); see also newer version [3]. URL: <http://eprint.iacr.org/2007/472>.
3. Jean-Philippe Aumasson, Simon Fischer, Shahram Khazaei, Willi Meier, Christian Rechberger, *New features of Latin dances: analysis of Salsa, ChaCha, and Rumba*, in [18] (2007), 470–488; see also older version [2]. Citations in this document: §7.
4. Rana Barua, Tanja Lange (editors), *Progress in cryptology—INDOCRYPT 2006, 7th international conference on cryptology in India, Kolkata, India, December 11–13, 2006, proceedings*, Lecture Notes in Computer Science, 4329, Springer, 2006. ISBN 3-540-49767-6. See [13].
5. Daniel J. Bernstein, *Floating-point arithmetic and message authentication* (2004). URL: <http://cr.yp.to/hash127.html>. ID dabadd3095644704c5cbe9690ea3738e. Citations in this document: §9.
6. Daniel J. Bernstein, *The Poly1305-AES message-authentication code*, in [14] (2005), 32–49. URL: <http://cr.yp.to/papers.html#poly1305>. ID 0018d9551b5546d97c340e0dd8cb5750. Citations in this document: §9, §9, §9.
7. Daniel J. Bernstein, *Salsa20 specification* (2005). URL: <http://cr.yp.to/snuffle.html>. Citations in this document: §7, §8, §8.
8. Daniel J. Bernstein, *Curve25519: new Diffie-Hellman speed records*, in [26] (2006), 207–228. URL: <http://cr.yp.to/papers.html#curve25519>. ID 4230efdfa673480fc079449d90f322c0. Citations in this document: §2, §2, §2, §2, §2, §5.
9. Daniel J. Bernstein, *Polynomial evaluation and message authentication* (2007). URL: <http://cr.yp.to/papers.html#pema>. ID b1ef3f2d385a926123e1517392e20f8c. Citations in this document: §9.

10. Daniel J. Bernstein, *Extending the Salsa20 nonce* (2008). URL: <http://cr.yp.to/papers.html#xsalsa>. ID c4b172305ff16e1429a48d9434d50e8a. Citations in this document: §7, §7.
11. John Black, Martin Cochran, *MAC reforgeability*, to appear, FSE 2009 proceedings (2009). Citations in this document: §9.
12. Paul Crowley, *Truncated differential cryptanalysis of five rounds of Salsa20*, in Workshop Record of SASC 2006: Stream Ciphers Revisted, eSTREAM technical report 2005/073 (2005). URL: <http://www.ecrypt.eu.org/stream/papers.html>. Citations in this document: §7.
13. Simon Fischer, Willi Meier, Côme Berbain, Jean-François Biasse, Matthew J. B. Robshaw, *Non-randomness in eSTREAM candidates Salsa20 and TSC-4*, in [4] (2006), 2–16. Citations in this document: §7.
14. Henri Gilbert, Helena Handschuh (editors), *Fast software encryption: 12th international workshop, FSE 2005, Paris, France, February 21–23, 2005, revised selected papers*, Lecture Notes in Computer Science, 3557, Springer, 2005. ISBN 3–540–26541–4. See [6].
15. Helena Handschuh, Bart Preneel, *Key-recovery attacks on universal hash function based MAC algorithms*, in [22] (2008), 144–161. Citations in this document: §9.
16. David A. McGrew, Scott R. Fluhrer, *Multiple forgery attacks against Message Authentication Codes* (2005). URL: <http://eprint.iacr.org/2005/161>. Citations in this document: §9.
17. Victor S. Miller, *Use of elliptic curves in cryptography*, in [25] (1986), 417–426. MR 88b:68040. Citations in this document: §2.
18. Kaisa Nyberg (editor), *Fast software encryption, 15th international workshop, FSE 2008, Lausanne, Switzerland, February 10–13, 2008, revised selected papers* (2008). ISBN 978-3-540-71038-7. See [3].
19. Kaisa Nyberg, Henri Gilbert, Matt Robshaw, *Galois MAC with forgery probability close to ideal* (2005). URL: http://csrc.nist.gov/groups/ST/toolkit/BCM/documents/comments/General_Comments/papers/Nyberg_Gilbert_and_Robshaw.pdf. Citations in this document: §9.
20. William Stein (editor), *Sage Mathematics Software (Version 3.0.2)*, The Sage Group, 2008. URL: <http://www.sagemath.org>. Citations in this document: §1.
21. Yukiyasu Tsunoo, Teruo Saito, Hiroyasu Kubo, Tomoyasu Suzaki, Hiroki Nakashima, *Differential Cryptanalysis of Salsa20/8*, in Workshop Record of SASC 2007: The State of the Art of Stream Ciphers, eSTREAM report 2007/010 (2007). URL: <http://www.ecrypt.eu.org/stream/papers.html>. Citations in this document: §7.
22. David Wagner (editor), *Advances in Cryptology—CRYPTO 2008, 28th annual international cryptology conference, Santa Barbara, CA, USA, August 17–21, 2008, proceedings*, Lecture Notes in Computer Science, 5157, Springer, 2008. ISBN 978-3-540-85173-8. See [15].
23. Mark N. Wegman, J. Lawrence Carter, *New classes and applications of hash functions*, in [1] (1979), 175–182; see also newer version [24]. URL: <http://cr.yp.to/bib/entries.html#1979/wegman>.
24. 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 [23]. ISSN 0022–0000. MR 82i:68017. URL: <http://cr.yp.to/bib/entries.html#1981/wegman>. Citations in this document: §9.
25. Hugh C. Williams (editor), *Advances in cryptology: CRYPTO '85*, Lecture Notes in Computer Science, 218, Springer, Berlin, 1986. ISBN 3–540–16463–4. See [17].

26. Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, Tal Malkin (editors), *9th international conference on theory and practice in public-key cryptography, New York, NY, USA, April 24–26, 2006, proceedings*, Lecture Notes in Computer Science, 3958, Springer, Berlin, 2006. ISBN 978–3–540–33851–2. See [8].