

## On the (in)security of ElGamal in OpenPGP

Luca De Feo, Bertram Poettering and Alessandro Sorniotti IBM Research Zürich

April 14, 2022, Real World Crypto, Amsterdam

## **Details matter**

## How to hang a picture? (ISO 3103<sup>1</sup>/<sub>2</sub>)



De Feo, Poettering, Sorniotti (IBM Research)

## **Details matter**

# How to hang a picture? (ISO 3103<sup>1</sup>/<sub>2</sub>)

1. Take hammer,

De Feo, Poettering, Sorniotti (IBM Research)

## **Details matter**

# How to hang a picture? (ISO 3103<sup>1</sup>/<sub>2</sub>)

- 1. Take hammer,
- 2. Strike nail...

Cryptographic standards, what's the worse that could happen?

- Theoretical break.
- Side-channel leakage.
- Implementations secure in isolation, do not interoperate.

## • Implementations secure in isolation, insecure when interoperating.

## **OpenPGP**

- An IETF Encryption standard, since 1998.
- One of two standards for end-to-end email encryption (along with S/MIME).
- Many implementations:

GnuPG, Botan (rnp/Thunderbird), Go (Protonmail), Libcrypto++, ...

## IETF RFCs:

RFC 4880 **OpenPGP Message Format**RFC 3156 MIME Security with OpenPGP
RFC 5581 The Camellia Cipher in OpenPGP
RFC 6637 Elliptic Curve Cryptography in OpenPGP

4/17

## **OpenPGP** algorithms

## Hash Functions: MD5, RIPE-MD, SHA-1, SHA-2. Symmetric Ciphers: IDEA, TripleDES, CAST5, Blowfish, AES, Twofish, Camellia. Public Key Encryption: RSA, ElGamal, ECDH. Signature Algorithms: RSA, DSA, ECDSA.

#### RFC 4880 (dated November 2007)

*"Implementations MUST implement DSA for signatures, and ElGamal for encryption. <i>Implementations SHOULD implement RSA* [...]*"* 

Public key algorithms specifications in OpenPGP

RSA PKCS #1

ECDH NIST SP 800-56A + RFC 6637

**DSA** FIPS 186-2

ECDSA FIPS 186-3

ElGamal El Gamal '85 / Handbook of Applied Cryptography '97

## ElGamal according to the OpenPGP standard?

#### 8.4.1 Basic ElGamal encryption

8.17 Algorithm Key generation for ElGamal public-key encryption

SUMMARY: each entity creates a public key and a corresponding private key. Each entity A should do the following:

- 1. Generate a large random prime p and a generator  $\alpha$  of the multiplicative group  $\mathbb{Z}_{n}^{*}$  of the integers modulo p (using Algorithm 4.84).
- 2. Select a random integer  $a, 1 \leq a \leq p-2$ , and compute  $\alpha^a \mod p$  (using Algorithm 2.143).
- 3. A's public key is  $(p, \alpha, \alpha^a)$ : A's private key is a.

©1997 by CRC Press, Inc. — See accompanying notice at front of chapter.

#### 8.18 Algorithm ElGamal public-key encryption

SUMMARY: B encrypts a message m for A, which A decrypts.

1. Encryption. B should do the following:

(a) Obtain A's authentic public key  $(p, \alpha, \alpha^a)$ .

- (b) Represent the message as an integer m in the range  $\{0, 1, \dots, p-1\}$ .
- (c) Select a random integer  $k, 1 \le k \le p-2$ .
- (d) Compute  $\gamma = \alpha^k \mod p$  and  $\delta = m \cdot (\alpha^a)^k \mod p$ .
- (e) Send the ciphertext  $c = (\gamma, \delta)$  to A.
- 2. Decryption. To recover plaintext m from c, A should do the following:
  - (a) Use the private key a to compute  $\gamma^{p-1-a} \mod p$  (note:  $\gamma^{p-1-a} = \gamma^{-a} = \gamma^{-a}$  $\alpha^{-ak}$ ).

(b) Recover m by computing  $(\gamma^{-a}) \cdot \delta \mod p$ .

#### II. THE PUBLIC KEY SYSTEM

First, the Diffie-Hellman key distribution scheme is reviewed. Suppose that A and B want to share a secret  $K_{AB}$ , where A has a secret x, and B has a secret  $x_{B}$ . Let p be a large prime and  $\alpha$  be a primitive element mod p, both known. A computes  $y_4 \equiv \alpha^{x_4} \mod p$ , and sends  $y_4$ . Similarly, B computes  $y_{B} \equiv \alpha^{x_{B}} \mod p$  and sends  $y_{B}$ . Then the secret  $K_{AB}$  is computed as

$$\begin{split} K_{AB} &\equiv \alpha^{x_A x_B} \bmod p \\ &\equiv y_A^{x_B} \bmod p \\ &\equiv y_B^{x_A} \bmod p. \end{split}$$

M This

Th

295

In any of the cryptographic systems based on discrete logarithms, p must be chosen such that p - 1 has at least one large prime factor. If p-1 has only small prime Gran Cryp factors, then computing discrete logarithms is easy (see [8]). Now suppose that A wants to send B a message m, Univ 1501 where 0 < m < p - 1. First A chooses a number k uniformly between 0 and p - 1. Note that k will serve as the secret  $x_{A}$  in the key distribution scheme. Then A computes the "key"

$$K \equiv y_B^k \mod p, \tag{1}$$

where  $v_n \equiv \alpha^{x_B} \mod p$  is either in a public file or is sent by B. The encrypted message (or ciphertext) is then the pair  $(c_1, c_2)$ , where

$$c_1 \equiv \alpha^k \mod p \qquad c_2 \equiv Km \mod p \qquad (2$$

and K is computed in (1).

ElGamal in the wild (OpenPGP ecosystem)

Large prime pSafe prime "Schnorr" prime "Lim-Lee" prime other Generator  $\alpha$ primitive element generates subgroup "short exponent" optimisation Private key 0 < a < pEphemeral key 0 < k < p"short exponent" optimisation

8/17

# What could possibly go wrong?



# Mars Climate Orbiter



Mars Climate Orbiter spacecraft

#### credit: Wikipedia

Apr 14, 2022, RWC Amsterdam 9 / 17

- Each of GnuPG, Botan and Libcrypto++ implements ElGamal in a different, non-RFC-4880-compliant way:
  - Each is secure taken in isolation.
  - They are interoperable: functionally and securely.

- Each of GnuPG, Botan and Libcrypto++ implements ElGamal in a different, non-RFC-4880-compliant way:
  - Each is secure taken in isolation.
  - They are interoperable: functionally and securely.

### • Go does not implement ElGamal key generation and is the least offender.

- Each of GnuPG, Botan and Libcrypto++ implements ElGamal in a different, non-RFC-4880-compliant way:
  - Each is secure taken in isolation.
  - They are interoperable: functionally and securely.
- We analyse 800K registered PGP ElGamal public keys:
  - 2K of them are exposed to practical plaintext recovery when GnuPG, Botan, Libcrypto++ (or any other library using the "short exponent" optimisation) encrypts to them. We call these cross-configuration attacks.
- Go does not implement ElGamal key generation and is the least offender.

- Each of GnuPG, Botan and Libcrypto++ implements ElGamal in a different, non-RFC-4880-compliant way:
  - Each is secure taken in isolation.
  - They are interoperable: functionally and securely.
- We analyse 800K registered PGP ElGamal public keys:
  - 2K of them are exposed to practical plaintext recovery when GnuPG, Botan, Libcrypto++ (or any other library using the "short exponent" optimisation) encrypts to them. We call these cross-configuration attacks.
- Go does not implement ElGamal key generation and is the least offender.
- We find side channels leaking ElGamal secret keys in GnuPG, Go and Libcrypto++:
  - GnuPG claimed to be side-channel resistant.
  - Our attack against GnuPG becomes more powerful in the cross-configuration scenario.

## **Prime generation**

Goal: prime p with at least one large prime factor q|(p-1).

Safe primes: p = 2q + 1:

• Considered kind of expensive, back in the '90s.

"Lim-Lee" primes:  $p = 2q_1q_2\cdots q_r$ , all  $q_i$  large:

- Cheaper than safe primes,
- Protecting against the same attacks.

"Schnorr" primes: p = 2qf + 1, with f arbitrary:

- Cheapest,
- Popularized by Schnorr signatures, DSA, FIPS-186-2.

Random primes: risky, don't do it!

Other: your imagination is the only limitation!

De Feo, Poettering, Sorniotti (IBM Research)

On the (in)security of ElGamal in OpenPGI



De Feo, Poettering, Sorniotti (IBM Research)



De Feo, Poettering, Sorniotti (IBM Research)

# Discrete log: when $\alpha$ is primitive $p-1 = 2 \cdot q \cdot \ell_3 \cdot \ell_4 \cdots$

 $lpha^x$ 







De Feo, Poettering, Sorniotti (IBM Research)

On the (in)security of ElGamal in OpenPG

pprimeα mod pgenerator

 $\alpha^x = X$  public key

p	prime	m	message
$lpha \mod p$	generator	y	random
$lpha^x = X$	public key		

p	prime	m	message
$lpha \mod p$	generator	y	random
$\alpha^x = X$	public key		

$$egin{array}{lll} (Y=lpha^y, & X^y\cdot m) & ext{encryption} \ & m=X^y\cdot m/Y^x & ext{decryption} \end{array}$$

# ElGamal Encryption p prime

 $lpha \mod p$  generator y random

 $\alpha^x = X$  public key

 $(p-1)= egin{array}{cccc} {
m safe} & {
m Schnorr} & {
m Lim-Lee} \ 2\cdot q & 2\cdot f\cdot q & 2\cdot q\cdot q_2\cdots q_r \end{array}$   $(q ext{ are "large" primes})$ 

- lpha generates all of  $\mathbb{Z}_p^*$  generates subgroup of order q (other possible)
  - $x \in [1, p-1]$  "short"

 $y \in [1, p-1]$  "short"

message

m



Bingo 0:key recovery from public key only (van Oorschot–Wiener)Bingo 1:message recovery from single ciphertext (this work)



GnuPG: Lim-Lee, generates all  $\mathbb{Z}_p^*$ , short exponents.



GnuPG: Lim-Lee, generates all  $\mathbb{Z}_p^*$ , short exponents.

Libcrypto++/Botan: safe primes, generates subgroup, short exponents.



GnuPG: Lim-Lee, generates all  $\mathbb{Z}_p^*$ , short exponents.

Libcrypto++/Botan: safe primes, generates subgroup, short exponents.

Go: no key generation,  $y \in [1, p - 1]$ .



GnuPG: Lim-Lee, generates all  $\mathbb{Z}_p^*$ , short exponents.

Libcrypto++/Botan: safe primes, generates subgroup, short exponents.

Go: no key generation,  $y \in [1, p - 1]$ .



## 800K registered OpenPGP ElGamal public keys

prime type	group size		quantity		
	p - 1	q	other	total	since 2016
Safe prime I	х			472,518	783
Safe prime II		Х		107,339	219
Lim–Lee I	?			211,271	6,003
Lim–Lee II		?		47	24
Quasi-safe I	х			15,592	89
Quasi-safe II		Х		20	3
Quasi-safe III			х	26,199	125
Schnorr I	?			828	810
Schnorr II		?		27	26
Schnorr III			Х	1,304	1,300

## Side channel vulnerabilities in exponentiation $\rightarrow$ Key recovery

## Threat model • Co-located attacker;

- Targets the exponentiation in the decryption routine;
- Must trigger decryption (e.g., email decryption).

Techniques FLUSH+RELOAD (instruction cache), PRIME+PROBE (data cache).

Findings



\*Verified experimentally on 2048 bits key.

- Each of GnuPG, Botan and Libcrypto++ implements ElGamal in a different, non-RFC-4880-compliant way:
  - Each is secure taken in isolation.
  - They are interoperable: functionally and securely.
- We analyse 800K registered PGP ElGamal public keys:
  - 2K of them are exposed to practical plaintext recovery when GnuPG, Botan, Libcrypto++ (or any other library using the "short exponent" optimisation) encrypts to them. We call these cross-configuration attacks.
- Go does not implement ElGamal key generation and is the least offender.
- We find side channels leaking ElGamal secret keys in GnuPG, Go and Libcrypto++:
  - GnuPG claimed to be side-channel resistant.
  - Our attack against GnuPG becomes more powerful in the cross-configuration scenario.