

Practical realisation and elimination of an ECC-related software bug attack^{*}

B. B. Brumley¹, M. Barbosa², D. Page³, and F. Vercauteren⁴

¹ Department of Information and Computer Science,
Aalto University School of Science, P.O. Box 15400, FI-00076 Aalto, Finland.

`billy.brumley@aalto.fi`

² HASLab/INESC TEC

Universidade do Minho, Braga, Portugal.

`mbb@di.uminho.pt`

³ Department of Computer Science, University of Bristol,
Merchant Venturers Building, Woodland Road, Bristol, BS8 1UB, UK.

`page@cs.bris.ac.uk`

⁴ Department of Electrical Engineering, Katholieke Universiteit Leuven,
Kasteelpark Arenberg 10, B-3001 Leuven-Heverlee, Belgium.

`fvercaut@esat.kuleuven.ac.be`

Abstract. We analyse and exploit implementation features in OpenSSL version 0.9.8g which permit an attack against ECDH-based functionality. The attack, although more general, can recover the entire (static) private key from an associated SSL server via 633 adaptive queries when the NIST curve P-256 is used. One can view it as a software-oriented analogue of the bug attack concept due to Biham et al. and, consequently, as the first bug attack to be successfully applied against a real-world system. In addition to the attack and a posteriori countermeasures, we show that formal verification, while rarely used at present, is a viable means of detecting the features which the attack hinges on. Based on the security implications of the attack and the extra justification posed by the possibility of intentionally incorrect implementations in collaborative software development, we conclude that applying and extending the coverage of formal verification to augment existing test strategies for OpenSSL-like software should be deemed a worthwhile, long-term challenge.

Keywords: elliptic curve, OpenSSL, NIST, fault attack, bug attack.

1 Introduction

Concrete implementation of cryptographic primitives is becoming easier as a result of more mature techniques and literature. Elliptic Curve Cryptography (ECC) is a case in point: twenty years ago ECC was limited to experts, but is now routinely taught in undergraduate engineering courses. However, such implementation tasks are still hugely challenging. This is because as well as

^{*} This is the full version of a shorter paper to appear at CT-RSA 2012.

functional correctness, the quality of an implementation is, in part, dictated by efficiency (e.g., execution speed and memory footprint) and physical security.

For (at least) two reasons, the efficiency of cryptographic primitives is an important issue within many applications. On one hand, many primitives represent an inherently expensive workload comprised of computationally-bound, highly numeric kernels. On the other hand, said primitives are often required in high-volume or high-throughput applications; examples include encryption of VPN traffic and full-disk encryption, both of which represent vital components in e-business. Both reasons are amplified because the primitive in question will often represent pure overhead at the application level. That is, cryptography is often an implicit enabling technology rather than an explicit feature: there is evidence to show it is common (and perhaps sane [17]) for users to disable security features in an application if it improves performance or responsiveness.

To summarise, some engineer must find an efficient way to map a complex, high-level specification of some primitive onto the characteristics of a demanding target platform, potentially using low-level programming languages and tools. Both the semantic gap between specification and implementation, and the skills gap between cryptography and engineering can be problematic. Two examples of the problems encountered, both relating to components in modern e-business work-flows, are as follows:

1. Nguyen [28] described an attack on GPG version 1.2.3, an open-source implementation of the OpenPGP standard. In short, the size of some security-critical parameters had been reduced; this meant computation was faster, but that the system as a whole was vulnerable to attack.
2. In part because of such wide-spread use, the open-source OpenSSL library has been subject to numerous attacks. Examples include issues relating to random number generation⁵, and badly formulated control-flow logic allowing malformed signatures to be reported as valid⁶.

Although other factors clearly contribute, one could argue that overly zealous optimisation is a central theme in both cases. Focusing on the provision of ECC in OpenSSL version 0.9.8g, this paper presents further evidence along similar lines. We stress that our aim is not to implicitly or explicitly devalue OpenSSL: one can, and *should*, read the paper more as a case study on the difficulty of cryptographic software implementation.

At the crux is an arithmetic bug, initially reported on the `openssl-dev` mailing list [30] in 2007 and later traced to the modular arithmetic underlying implementation of specific NIST elliptic curves; in short, the bug causes modular multiplications to (transiently) produce incorrect output. To the best of our knowledge, no cryptanalytic exploitation of this bug was previously known. Perhaps for this reason, it has not been considered a security risk, but rather a minor issue of functionality. Indeed, although the bug has been resolved in OpenSSL versions 0.9.8h and later it persists⁷; for example versions of the li-

⁵ http://www.openssl.org/news/secadv_20071129.txt

⁶ http://www.openssl.org/news/secadv_20090107.txt

⁷ <http://marc.info/?t=131401133400002>

library are deployed in (at least) two major Linux distributions, namely Debian (as late as 5.0 “Lenny”) and Ubuntu (as late as 9.10 “Karmic”).

The main contribution of this paper is a concrete attack: we show how the bug can be exploited to mount a full key recovery attack against implementations of Elliptic Curve Diffie-Hellman (ECDH) key agreement. The nature of the bug means the attack represents a software analogue (or a first practical realisation) of the bug attack concept [8] due to Biham et. al. Our attack works whenever the ECDH public key is static, and therefore reused across several key agreement protocol executions. In particular, any higher-level application relying on the SSL/TLS implementation of `OpenSSL` in the following two scenarios could be vulnerable:

1. Use of *static* ECDH-based cipher suites, (e.g., ECDH-ECDSA and ECDH-RSA). In such cipher suites, the TLS server holds a public key certificate that directly authenticates the ECDH public key; this is shared across an arbitrary number of key exchanges.
2. Use of *ephemeral* ECDH-based cipher suites (e.g., ECDHE-ECDSA and ECDHE-RSA) in combination with the `OpenSSL` *ephemeral-static* ECDH optimisation. In such cipher suites, and according to the TLS specification, a fresh ECDH public key should be generated for each key exchange. However `OpenSSL` allows one-time generation of said key when the TLS server is initialised, sharing it across an arbitrary number of key exchanges thereafter.

As a concrete example, we demonstrate the attack on `stunnel` version 4.42 (when linked against `OpenSSL` version 0.9.8g), an SSL-based tunnelling proxy.

As well as discussing potential countermeasures for vulnerable versions of the library, we also explore an alternative, longer-term solution. Specifically, we investigate use of formal verification as a means to *prevent* similar bugs rather than just detecting them a posteriori. This approach is particularly relevant in addressing the possibility of *intentionally* incorrect implementations, which could constitute a serious risk in software components developed using an open, collaborative approach. Our conclusion is that although such techniques can already play an important role, a step-change in attitude toward their use is required as software complexity increases; despite the effort required to adopt a development strategy that supports formal verification, this seems an important area for future work in the context of `OpenSSL`-like software.

From here on we use `OpenSSL` as a synonym for `OpenSSL` version 0.9.8g unless otherwise stated. In Section 2 we present a detailed analysis of features in `OpenSSL` that permit our attack to work, before explaining the attack itself, and possible countermeasures, in Section 3. In Section 4 we discuss approaches to formal verification that could help prevent similar defects in `OpenSSL`, and therefore similar attacks, and offer some concluding remarks in Section 5.

2 Background and analysis

The aim of this section is to relate high-level, standardised ECC with an analysis of associated low-level implementation features in `OpenSSL` which support our attack.

2.1 Elliptic Curves

The use of elliptic curves in cryptography was pioneered during the mid 1980s by Miller [25] and Koblitz [22]; we refer the reader to, for example, Menezes et al. [16] for a more complete overview of the topic. Briefly, an elliptic curve E over a finite field K is defined by the general Weierstraß equation, for $a_i \in K$

$$E(K) : y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6.$$

The K -rational points on a curve E , i.e., those $(x, y) \in K^2$ which satisfy the curve equation, plus the point at infinity \mathcal{O} , form an additive group under a group law defined by the chord-tangent process.

Most ECC-based protocols use this additive group structure as a means to present a discrete logarithm problem as the basis for security. The Elliptic Curve Discrete Logarithm Problem (ECDLP) is constructed by considering scalar multiplication of some point $P \in E$ by an integer k , i.e., $Q = [k]P$.

2.2 NIST Standard Curves

To support efficient, interoperable implementations of ECC, NIST (among others) specify several standard parametrisations; once agreed, this provides a group \mathbb{G} in which a given protocol can operate. In the case where $K = \mathbb{F}_p$, i.e., the curve is defined over a finite field of large prime characteristic, there are basically three things to note:

1. The value of p is selected to have a special form, in particular to be a Generalised Mersenne (GM) prime [32]. This form affords very efficient modular reduction with no need for division or special representations (e.g., Montgomery representation [26]).
2. The curve equation is simplified to $E(\mathbb{F}_p) : y^2 = x^3 + a_4x + a_6$, where $a_4 = -3$. This acts to simplify arithmetic on E without negative consequence.
3. Various sets of domain parameters $D = \{p, a_4, a_6, x_G, y_G, n, h\}$ are specified. Each set of domain parameters provides $\mathbb{G} = \langle G \rangle$, a large prime elliptic curve subgroup of order n generated by $G = (x_G, y_G) \in E$; note that the cofactor h is such that $|E| = n \cdot h$ but that for all NIST curves $h = 1$.

2.3 OpenSSL Implementation of Point Validation

The following definition offers a mechanism to validate points; the aim is to prevent invalid curve attacks of the type described above:

```

315     /* testing whether the pub_key is on the elliptic curve */
316     if (!EC_POINT_is_on_curve(eckey->group, eckey->pub_key, ctx))
317     {
318         ECerr(EC_F_EC_KEY_CHECK_KEY, EC_R_POINT_IS_NOT_ON_CURVE);
319         goto err;
320     }
321     /* testing whether pub_key * order is the point at infinity */
322     if (!EC_GROUP_get_order(eckey->group, order, ctx))
323     {
324         ECerr(EC_F_EC_KEY_CHECK_KEY, EC_R_INVALID_GROUP_ORDER);
325         goto err;
326     }
327     if (!EC_POINT_copy(point, eckey->pub_key))
328     {
329         ECerr(EC_F_EC_KEY_CHECK_KEY, ERR_R_EC_LIB);
330         goto err;
331     }
332     if (!EC_POINT_mul(eckey->group, point, order, NULL, NULL, ctx))
333     {
334         ECerr(EC_F_EC_KEY_CHECK_KEY, ERR_R_EC_LIB);
335         goto err;
336     }
337     if (!EC_POINT_is_at_infinity(eckey->group, point))
338     {
339         ECerr(EC_F_EC_KEY_CHECK_KEY, EC_R_WRONG_ORDER);
340         goto err;
341     }

```

Fig. 1. A fragment of `EC_KEY_check_key` demonstrating explicit point validation in OpenSSL.

```

1338 /* rh := X^2 */
1339 if (!field_sqr(group, rh, &point->X, ctx)) goto err;
1340
1341 if (!point->Z_is_one)
1342 {
1343     if (!field_sqr(group, tmp, &point->Z, ctx)) goto err;
1344     if (!field_sqr(group, Z4, tmp, ctx)) goto err;
1345     if (!field_mul(group, Z6, Z4, tmp, ctx)) goto err;
1346
1347     /* rh := (rh + a*Z^4)*X */
1348     if (group->a_is_minus3)
1349     {
1350         if (!BN_mod_lshift1_quick(tmp, Z4, p)) goto err;
1351         if (!BN_mod_add_quick(tmp, tmp, Z4, p)) goto err;
1352         if (!BN_mod_sub_quick(rh, rh, tmp, p)) goto err;
1353         if (!field_mul(group, rh, rh, &point->X, ctx)) goto err;
1354     }
1355     else
1356     {
1357         if (!field_mul(group, tmp, Z4, &group->a, ctx)) goto err;
1358         if (!BN_mod_add_quick(rh, rh, tmp, p)) goto err;
1359         if (!field_mul(group, rh, rh, &point->X, ctx)) goto err;
1360     }
1361
1362     /* rh := rh + b*Z^6 */
1363     if (!field_mul(group, tmp, &group->b, Z6, ctx)) goto err;
1364     if (!BN_mod_add_quick(rh, rh, tmp, p)) goto err;
1365 }
1366 else
1367 {
1368     /* point->Z_is_one */
1369
1370     /* rh := (rh + a)*X */
1371     if (!BN_mod_add_quick(rh, rh, &group->a, p)) goto err;
1372     if (!field_mul(group, rh, rh, &point->X, ctx)) goto err;
1373     /* rh := rh + b */
1374     if (!BN_mod_add_quick(rh, rh, &group->b, p)) goto err;
1375 }
1376
1377 /* 'lh' := Y^2 */
1378 if (!field_sqr(group, tmp, &point->Y, ctx)) goto err;
1379
1380 ret = (0 == BN_ucmp(tmp, rh));

```

Fig. 2. A fragment of `ec_GFp_simple_is_on_curve`, called by `ec_GFp_simple_oct2point`, demonstrating implicit point validation in OpenSSL.

Definition 1. [4, Section 3] A point $Q = (x_Q, y_Q)$ (static or ephemeral) associated with a set of domain parameters D is valid if the following four conditions are satisfied:

1. $Q \neq \mathcal{O}$.
2. x_Q and y_Q are properly represented elements of \mathbb{F}_p (e.g., integers in the interval $[0, p - 1]$).
3. Q satisfies the defining equation of the elliptic curve E .
4. $[n]Q = \mathcal{O}$.

If any one of these conditions is violated, then Q is invalid.

Although condition 4 dominates the computation time, if $h = 1$ (as per NIST standard curves) then explicitly verifying said condition is redundant. Quoting from [4, Section 3]:

There may be ways of verifying condition 4 . . . that are much faster than performing an expensive point multiplication $[n]Q$. For example, if $h = 1$ (which is usually the case for elliptic curves over prime fields that are used in practice), then condition 4 is implied by the other three conditions.

As such, one can identify two subtly different methods of implementing point validation: *explicit validation* where the validity of conditions 1 to 4 are each explicitly verified regardless of h , and *Implicit validation* where the validity of conditions 1 to 3 are each explicitly verified, but since $h = 1$ condition 4 is implicitly verified.

To highlight the fact that *both* approaches are applied in different contexts, consider the following standards:

ANSI X9.62 [2] contains two relevant algorithms:

1. in 5.2.2, PUBLIC KEY VALIDATION (OPTIONAL) performs an explicit point validation (but this is not mandatory),
2. in 4.3.7, OCTET-STRING-TO-POINT CONVERSION contains an implicit point validation.

ANSI X9.63 [3] contains three relevant algorithms:

1. 5.2.2.1, STANDARD PUBLIC KEY VALIDATION PRIMITIVE performs an explicit point validation,
2. 5.2.2.2, EMBEDDED PUBLIC KEY VALIDATION PRIMITIVE performs an implicit point validation,
3. 4.3.7, OCTET-STRING-TO-POINT CONVERSION contains an implicit point validation.

P1363 [21] contains one explicit and one implicit point validation algorithm within A.16.10, ALGORITHMS FOR VALIDATING EC PUBLIC KEYS.

OpenSSL implements and uses ANSI OCTET-STRING-TO-POINT CONVERSION and contains a function `EC_KEY_check_key` that performs explicit point validation; interestingly, this only seems to be used internally for testing the library at compile-time.

2.4 OpenSSL Implementation of NIST Standard Curves

For a GM-prime p , multi-precision integer multiplication modulo p can be particularly efficient. OpenSSL uses this fact to support ECC implementations over the NIST standard curves P-192, P-224, P-256, P-384 and P-521. Using P-256 as an example, we have

$$p = 2^{256} - 2^{224} + 2^{192} + 2^{96} - 1$$

and, from here on, we refer to the resulting elliptic curve as E .

Assuming a processor with a 32-bit word size, imagine that given two 8-word operands $0 \leq x, y < p$, the goal is to compute $x \cdot y \pmod{p}$. Solinas demonstrates [32, Example 3, Page 20] that given $z = x \cdot y$, the 16-word integer product of x and y , one can compute $z \pmod{p}$ by first forming nine 8-word intermediate values

$$\begin{aligned} S_0 &= (z_7, z_6, z_5, z_4, z_3, z_2, z_1, z_0) \\ S_1 &= (z_{15}, z_{14}, z_{13}, z_{12}, z_{11}, 0, 0, 0) \\ S_2 &= (0, z_{15}, z_{14}, z_{13}, z_{12}, 0, 0, 0) \\ S_3 &= (z_{15}, z_{14}, 0, 0, 0, z_{10}, z_9, z_8) \\ S_4 &= (z_8, z_{13}, z_{15}, z_{14}, z_{13}, z_{11}, z_{10}, z_9) \\ S_5 &= (z_{10}, z_8, 0, 0, 0, z_{13}, z_{12}, z_{11}) \\ S_6 &= (z_{11}, z_9, 0, 0, z_{15}, z_{14}, z_{13}, z_{12}) \\ S_7 &= (z_{12}, 0, z_{10}, z_9, z_8, z_{15}, z_{14}, z_{13}) \\ S_8 &= (z_{13}, 0, z_{11}, z_{10}, z_9, 0, z_{15}, z_{14}) \end{aligned}$$

and then computing $S \pmod{p}$ with

$$S = S_0 + 2S_1 + 2S_2 + S_3 + S_4 - S_5 - S_6 - S_7 - S_8. \quad (1)$$

Note that $|S|$ cannot be much larger than p , meaning a small number of extra modular additions or subtractions, depending on the sign of S , would give the correct result.

OpenSSL adopts a similar approach for P-192, P-224 and P-521 but deviates for P-256 and P-384: we again use P-256 as an example, but note that the same problem exists for P-384. It proceeds using the following faulty algorithm: first it computes $t = S \pmod{2^{256}}$ and the correct carry c (which is positive or negative) such that

$$S = t + c \cdot 2^{256}.$$

Note that per the comment above, the carry has a small magnitude; by inspection it is loosely bounded by $-4 \leq c \leq 6$, which is used from here on wlog. The result is computed, potentially incorrectly, via two steps:

1. set $r' = (t - c \cdot p) \pmod{2^{256}}$, then
2. if $r' \geq p$, $r' = r' - p$.

The concrete implementation of these steps uses a fixed look-up table $T[i] = i \cdot p \pmod{2^{256}}$ for small i , by computing $r' = t - \text{sign}(c) \cdot T[|c|] \pmod{2^{256}}$. The modular reduction in this case is implicit, realised by truncating the result to 8 words.


```

485 nist_cp_bn_0(buf, a_d + BN_NIST_256_TOP, top - BN_NIST_256_TOP, BN_NIST_256_TOP);
486
487 /*S1*/
488 nist_set_256(t_d, buf, 15, 14, 13, 12, 11, 0, 0, 0);
489 /*S2*/
490 nist_set_256(t_d2, buf, 0, 15, 14, 13, 12, 0, 0, 0);
491 if (bn_add_words(t_d, t_d, t_d2, BN_NIST_256_TOP))
492     carry = 2;
493 /* left shift */
494 {
495     register BN_ULONG *ap, t, c;
496     ap = t_d;
497     c=0;
498     for (i = BN_NIST_256_TOP; i != 0; --i)
499     {
500         t = *ap;
501         *(ap++) = ((t << 1) | c) & BN_MASK2;
502         c = (t & BN_TBIT) ? 1 : 0;
503     }
504     if (c)
505         ++carry;
506 }
507
508 if (bn_add_words(r_d, r_d, t_d, BN_NIST_256_TOP))
509     ++carry;
510 /*S3*/
511 nist_set_256(t_d, buf, 15, 14, 0, 0, 0, 10, 9, 8);
512 if (bn_add_words(r_d, r_d, t_d, BN_NIST_256_TOP))
513     ++carry;
514 /*S4*/
515 nist_set_256(t_d, buf, 8, 13, 15, 14, 13, 11, 10, 9);
516 if (bn_add_words(r_d, r_d, t_d, BN_NIST_256_TOP))
517     ++carry;
518 /*D1*/
519 nist_set_256(t_d, buf, 10, 8, 0, 0, 0, 13, 12, 11);
520 if (bn_sub_words(r_d, r_d, t_d, BN_NIST_256_TOP))
521     --carry;
522 /*D2*/
523 nist_set_256(t_d, buf, 11, 9, 0, 0, 15, 14, 13, 12);
524 if (bn_sub_words(r_d, r_d, t_d, BN_NIST_256_TOP))
525     --carry;
526 /*D3*/
527 nist_set_256(t_d, buf, 12, 0, 10, 9, 8, 15, 14, 13);
528 if (bn_sub_words(r_d, r_d, t_d, BN_NIST_256_TOP))
529     --carry;
530 /*D4*/
531 nist_set_256(t_d, buf, 13, 0, 11, 10, 9, 0, 15, 14);
532 if (bn_sub_words(r_d, r_d, t_d, BN_NIST_256_TOP))
533     --carry;
534
535 if (carry)
536 {
537     if (carry > 0)
538         bn_sub_words(r_d, r_d, _256_data + BN_NIST_256_TOP *
539             --carry, BN_NIST_256_TOP);
540     else
541     {
542         carry = -carry;
543         bn_add_words(r_d, r_d, _256_data + BN_NIST_256_TOP *
544             --carry, BN_NIST_256_TOP);
545     }
546 }
547
548 r->top = BN_NIST_256_TOP;
549 bn_correct_top(r);
550 if (BN_ucmp(r, field) >= 0)
551 {
552     bn_sub_words(r_d, r_d, _nist_p_256, BN_NIST_256_TOP);
553     bn_correct_top(r);
554 }
555 bn_check_top(r);

```

Fig. 3. A program fragment demonstrating the faulty OpenSSL modular reduction algorithm for P-256.

The intention is to eliminate any possibility of overflow; the assumption is that c is the exact quotient of division of S by p .

The reasoning behind the faulty algorithm is that if one writes $S = t + c \cdot 2^{256}$, then the exact quotient $q = S \div p$ is given by

1. if $c \geq 0$, then $q = c$ or $q = c + 1$,
2. if $c < 0$, then $q = c$ or $q = c - 1$

since c is small. Indeed, write $\Delta = 2^{256} - p$, then after subtracting $c \cdot p$ we obtain

$$S - c \cdot p = t + c \cdot 2^{256} - c \cdot p = t + c \cdot \Delta.$$

Since $-4 \leq c \leq 6$ and $\Delta < 2^{224}$, this shows the result is bounded by $-p < t + c \cdot \Delta < 2p$. The faulty algorithm therefore computes an incorrect result in the following cases:

- If $c \geq 0$, the algorithm fails when $t + c \cdot \Delta \geq 2^{256}$ since it computes r' only modulo 2^{256} and not as a full integer (for which the resulting algorithm would have been correct). Note that in this case the correct result would be $r' + \Delta$ and that modulo p , the correct result thus is $r' + 2^{256} \pmod{p}$.
- If $c < 0$, the algorithm fails when $t + c \cdot \Delta < 0$. The correct result then depends on whether $(t + c \cdot \Delta) \pmod{2^{256}} \geq p$ or not: in the former case, the correct result is $r' - \Delta$, whereas in the latter case, the correct result is given by $r' + 2^{256} - 2\Delta$. Note that although there are two different subcases for $c < 0$, the errors $-\Delta$ and $2^{256} - 2\Delta$ are congruent modulo p , i.e. modulo p , the correct result is given by $r' - 2^{256} \pmod{p}$.

Note that Ciet and Joye [9, Section 3.2] consider the case of faults in the underlying field; the fault (resp. bug) here is certainly related, but occurs as a result of erroneous computation rather than corrupted parameters.

The resulting bug is difficult to detect using the (random) test vector approach employed by OpenSSL: it simply manifests itself too rarely. An upper bound for the probability of the bug being triggered can be obtained by ignoring the probabilities of certain carries occurring and analysing the case $t + 6 \cdot \Delta \geq 2^{256}$: if t was chosen uniformly over the interval $[0, 2^{256}[$, then this case occurs with probability less than 2^{-29} , so OpenSSL computes the result incorrectly with probability less than $10 \cdot 2^{-29}$.

To *deliberately* trigger the bug, we empirically arrived at the following strategies (after inspection of partial products within the integer product of x and y):

- For modular multiplication, selecting x, y as follows should induce an incorrect result for any random $0 \leq R_x, R_y < 2^{31}$:

$$\begin{aligned} x &= (2^{32} - 1) \cdot 2^{224} + 3 \cdot 2^{128} + R_x \\ y &= (2^{32} - 1) \cdot 2^{224} + 1 \cdot 2^{96} + R_y \end{aligned}$$

- For modular squaring, selecting $x = (2^{32} - 1) \cdot 2^{224} + 1 \cdot 2^{129} + R_x$, should induce an incorrect result for any random $0 \leq R_x < 2^{31}$.




Fig. 4. A description of ECDH key exchange.

2.5 ECC Cipher Suites for TLS

Modern versions⁸ of the Transport Layer Security (TLS) standard provide a number of different cipher suites that rely on ECC for key exchange. We focus on Elliptic Curve Diffie-Hellman (ECDH) and Ephemeral Elliptic Curve Diffie-Hellman (ECDHE) based cipher suites⁹. To be precise, in these cipher suites the key exchange protocol is conducted to establish a secret key for a session i between a client C and a server S ; it proceeds in three stages outlined in Figure 4, with client and server assumed to share $D = \{p, A, B, x_G, y_G, n, h\}$, a set of domain parameters. After the protocol terminates, $R_S^i = R_C^i$ represents the shared key.

Figure 5 illustrates the TLS handshake at a higher level of abstraction. We now describe how said handshake proceeds, detailing how the ECDH protocol messages formalised above are embedded in the communication. While our attacks are equally applicable in the case of client authentication, we omit the details for this case. The `ClientHello` message conveys the protocol version, supported cipher and compression methods, and a nonce to ensure freshness. The `ServerHello` message is analogous, but selects parameters from the methods proposed by the client (contingent on the server supporting them). The content of the `Certificate` message varies depending on the selected cipher suite:

- In ECDH-ECDSA, the `Certificate` message contains a static ECDH public key authenticated by a public key certificate signed with ECDSA; ECDH-RSA is analogous, but the public-key certificate is signed using an RSA signature. The static ECDH public key corresponds to the value Q_S^i above, and the server will reuse this value for multiple key exchanges with an arbitrary number of clients. For this reason, the `ServerKeyExchange` message is omitted in ECDH suites.
- In ECDHE-ECDSA, the `Certificate` message contains an ECDSA verification key, which is authenticated by a certificate signed with the same

⁸ <http://tools.ietf.org/html/rfc5246>

⁹ <http://tools.ietf.org/html/rfc4492>




Fig. 5. Message flow in a TLS handshake with ECDH cipher suites; messages relating to client authentication are omitted for clarity, and those marked with † are only sent under specific circumstances.

algorithm; ECDHE-RSA is analogous but an RSA signature verification key is sent, and the public-key certificate is signed using an RSA signature. The server also sends message **ServerKeyExchange**, containing both a fresh ephemeral ECDHE public key (i.e., Q_S^i) and a digital signature authenticating this and other handshake parameters, including the exchanged nonces. Said signature is produced with either the ECDSA or RSA signing key matching to the verification key sent in the **Certificate** message.

The **ServerHelloDone** message marks the end of this stage of the protocol; the client then sends its ephemeral ECDHE key in the **ClientKeyExchange** message, which always includes a fresh ephemeral Q_C^i . Finally, the negotiated symmetric cipher suite is activated via the **ChangeCipherSpec** message. A session key for the cipher suite is derived from $R_S^i = R_C^i$ and the exchanged nonces. The **Finished** messages provide key confirmation for both parties, notably occurring client-first, and depend on all previous messages in the protocol execution.

2.6 OpenSSL Implementation of the ECC Cipher Suites

The ECDH implementation in OpenSSL is seemingly straightforward, and follows the TLS specification. However, the ECDHE implementation offers two distinct options for server applications. The first follows the specification and generates a new ephemeral ECDH key pair for every protocol execution. Deviating from

the specification, the second features an optimisation termed ephemeral-static ECDH¹⁰.

When activated, the optimisation means a single ECDH key pair is generated during initialisation of the OpenSSL context within an application. This key pair is reused for all protocol executions thereafter; with ephemeral-static ECDH, OpenSSL has the server use a static key (i.e., a fixed k_i^S and hence Q_i^S for all i) for each OpenSSL context. Put another way, the key pair is ephemeral for each application instance and not (necessarily) per handshake instance. While this preserves forward secrecy between application instances, it violates forward secrecy within a single application instance when performing more than a single protocol execution. Interestingly, the default behaviour is the latter: to “opt out” and disable the optimisation, the application must explicitly use the `SSL_OP_SINGLE_ECDH_USE` option during initialisation of the context.

3 An Attack on ECDH in OpenSSL

Implementing scalar multiplication. For scalar multiplication on $E(\mathbb{F}_p)$, OpenSSL uses a textbook double-and-add algorithm along with the modified width- w NAF representation of k . For P-256 OpenSSL sets $w = 4$, i.e., each non-zero digit from digit set $\mathcal{D} = \{0, \pm 1, \pm 3, \pm 5, \pm 7\}$ is followed by at least three zero digits. Modified NAF is otherwise identical to traditional NAF but allows the most-significant digit to violate the non-adjacency property, if doing so does not increase the weight but reduces the length. This slight distinction between the two affects neither the derivation of our attack nor the implementation of it: henceforth we use NAF synonymously with traditional NAF.

Attack Goals and Limitations. The goal of the attacker is to recover the fixed k_S in the server-side computation of $R_S^i = [k_S]Q_C^i$. The algorithm we propose and implement recovers the digits of k_S by detecting faults in the server-side computation of R_S^i . The ability of the attacker to observe these faults heavily depends on the protocol and/or cryptosystem under attack. For example, when k_S is fixed but Q_C^i is not, it is uncommon for a protocol to directly reveal R_S^i to another participant. Inspecting TLS, one way the attacker can detect faults is by attempting to complete the handshake. If R_S^i is fault-free (denoted $R_S^i \in E$) then the computed session key is the same for both the client and server. Consider the phase of the handshake when the client sends `ChangeCipherSpec`, signalling that it has activated the newly negotiated session cipher suite, and transmits the encrypted `Finished` handshake message for key confirmation. Then, if the server successfully decrypts said message, the handshake continues, and ultimately succeeds. On the other hand, if R_S^i is faulty (denoted $R_S^i \notin E$) then the session keys differ, the server will not obtain the expected key confirmation message, and the handshake ultimately fails. The practical consequence is that the attacker *cannot* arbitrarily choose each Q_C^i in the protocol, rather he *must* know the discrete logarithm of said point in order to correctly calculate the negotiated session key.

¹⁰ <http://tools.ietf.org/html/rfc5753>

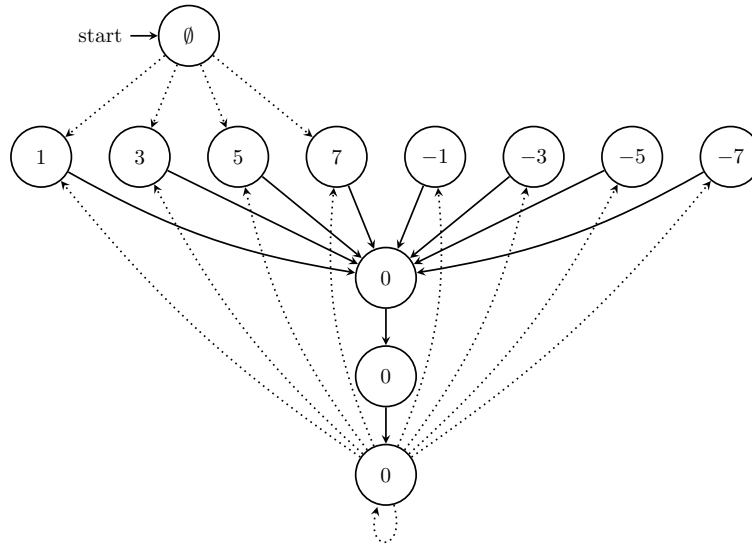


Fig. 6. The search space, illustrated: an automaton for constructing NAF strings.

The Attack Algorithm. Having established a method to detect the occurrence of faults, the algorithm proceeds in an exhaustive depth-first search for $\text{NAF}(k_S)$ starting from the most-significant digit, trimming limbs and backtracking by iteratively observing handshake results. Figure 6 contains an automaton for constructing NAF strings that illustrates the search space, starting from the empty string. At each node in the search, the attacker submits a different carefully chosen point, termed a *distinguisher point*, to determine if the next unknown digit takes a specific value at the given iteration, tracing the execution path of the server-side scalar multiplication. We define a distinguisher point for an integer prefix a and target digit $b \in \mathcal{D} \setminus \{0\}$ to be a point $D_{a,b} = [l]G \in E$ such that $[a \parallel b \parallel d]D_{a,b} \notin E$ and $[a \parallel c \parallel d]D_{a,b} \in E$ for all $c \in \mathcal{D} \setminus \{0, b\}$ both hold. Here, l is known, a is the known portion of $\text{NAF}(k_S)$, and d is any sufficiently long random padding string that completes the resulting concatenation to a valid NAF string. In practice, testing a single distinguisher point requires varying d over many values to ensure the computation reaches a sufficient number of possible subsequent algorithm states: this acts to deter false positives.

We step through the first few iterations to demonstrate the algorithm. For clarity, we use subscripts on variables a and \mathcal{D} to identify the iteration, i.e., the digit index in the NAF representation from most- to least-significant, we are referring to. For $i = 1$, a_1 is the empty string and $\mathcal{D}_1 = \{1, 3, 5, 7\}$. The attacker finds a distinguisher point $D_{\emptyset,b}$ for each $b \in \mathcal{D}_1 \setminus \{1\}$ and uses these three points

in attempted handshakes to the server¹¹. Handshake failure reveals the correct digit, and allows us to set a for the next iteration as $a_2 = b$; if all handshakes succeed, the attacker deduces $a_2 = 1$ for the next iteration. Enforcing NAF rules, for $i = 5$ we have $a_5 = a_2 \parallel 000$ and $\mathcal{D}_5 = \{0, \pm 1, \pm 3, \pm 5, \pm 7\}$. The attacker then finds $D_{a_5, b}$ for each $b \in \mathcal{D}_5 \setminus \{0\}$ and uses these eight points in attempted handshakes to the server. Handshake failure reveals the correct $a_6 = a_5 \parallel b$ and if all handshakes succeed the attacker deduces $a_6 = a_5 \parallel 0$. The attack continues in this manner to recover all subsequent digits. On average, our attack takes 4 handshake attempts to recover subsequent non-zero digits, and 8 handshake attempts to detect zero digits (note that we do not explicitly check for zeros which are implied by the NAF representation).

Relation to the Bug Attacks of Biham et al. This algorithm relates to that which Biham et al. used to mount a bug attack against Pohlig-Hellman in the \mathbb{F}_p^* setting [8, Section 4.1.1]. The authors consider recovering the binary digits of the exponent from a left-to-right binary modular exponentiation algorithm. It does this by finding input $X \in \mathbb{F}_p^*$ such that $(X^2)^2$ fails yet $(X^2)X$ does not, i.e., it uses the former to query explicitly for any zero digits and implicitly obtain any non-zero digits. Assume the attacker knows l such that $X = g^l$ (this is not necessary to carry out their attacks, but is necessary when adapting their strategy to attack TLS). The most-significant digit of the victim’s binary string is a non-zero digit by definition. The attacker queries the next digit by submitting X . Assume wlog. that it fails: the attacker now knows the two most-significant digits are 1 and 0. To obtain the (or analogy of a) distinguisher point for the next iteration, the attacker simply computes $X^{1/2}$ and, knowing the discrete logarithm of X to the base g , can easily derive the logarithm of this group element. This procedure essentially cancels out the effect of the known digits, forcing the accumulator to take value X at the intended iteration.

In the discussion that follows we will see that, in our attack, we search for all distinguishing points independently. Indeed, although it is tempting to use existing distinguisher points to derive subsequent points, which would reduce the complexity of our attack, this approach does not seem to apply in our scenario. The distinguisher point derivation technique by Biham et al. works for \mathbb{F}_p^* (and even $E(\mathbb{F}_p)$ when using affine coordinates) because there is a unique representation for group elements in both cases (i.e., elements of \mathbb{F}_p are stored as their smallest non-negative residue). There are a number of reasons why this strategy is ineffective for the OpenSSL implementation of ECC, the most prominent being the use of projective coordinates: group element representations are not unique, as the implementation computes in an equivalence class. As a result, the attacker cannot force the accumulator to the desired value since it undergoes projective point doublings. To summarise, there is no obvious way to cancel out

¹¹ When $i = 1$ finding a distinguisher point $D_{0,1}$ is less practical as it can cause the table of pre-computed points to be erroneously populated, so in this case querying for that particular digit value occurs implicitly.




Fig. 7. Left: distribution of required search steps to find distinguisher points (lg-linear). The solid line represents effort towards the beginning of the attack (mean 26.5 s.d. 2.0) and the dashed line towards the end (mean 26.9 s.d. 1.9). Right: distribution of required queries to the server, or distinguisher points, for the full attack (mean 633.2 s.d. 57.7).

the effect of the known digits. In any case, we will see in the following that this does not impact on the practicality of our attack in a significant way.

Finding Distinguisher Points. Lacking an analytical method to derive distinguisher points, our implementation of the attack resorts to random search. This is done by choosing l at random and testing whether $D_{a,b}$ satisfies the properties of a distinguisher point and, as previously mentioned, varying d over a reasonable amount of random values. The practicality of carrying out the attack thus hinges (in part) on the effort required to find said points, i.e., the average number of l values to test. Figure 7 (left) depicts our observed effort in two different parts of the attack. The solid line represents effort required at the beginning of the attack with less than 12 digits recovered. The dashed line represents effort required towards the end of the attack with roughly 224 digits recovered. This empirical data suggests not only that the computational effort to find a specific distinguisher point is rather modest, but that said effort does not significantly depend on the amount of known key digits. That is, as the attacker recovers successive digits, the effort to find new distinguisher points remains fairly constant. It is worth mentioning that the search scales perfectly with the number of computing nodes.

Attack Analysis and Results. The practicality of the attack additionally hinges on the number of required distinguisher points, i.e., the average number of attempted handshakes to recover the entire key. Our theoretical analysis of the expected number of queries, based on a rough approximation to the distribution of digits in the NAF representation of a random 256-bit secret exponent, points to an approximate value of 635 handshakes, suggesting that said value is similarly modest. We also measured this value empirically and obtained a consistent result, as illustrated in Figure 7 (right).

The proof-of-concept source code for our attack implementation includes distinguisher points for all NAF strings up to length 12. As it stands this immediately removes roughly 12 bits of entropy from the private key, and is of course easily extendable. The code includes instructions for running the attack in two different use cases:

1. The `stunnel` application provides a flexible SSL proxy for any application that does not natively support SSL; it links against `OpenSSL` to provide the SSL functionality. When `stunnel` is configured to support ECDH suites with a static ECDH key, our attack implementation facilitates recovery of said private key. Once carried out, it allows the attacker to decrypt all previous SSL sessions and to impersonate the server indefinitely.
2. The `s.server` application within `OpenSSL` is a generic SSL server. This application, and those similar to it supporting ECDHE suites, are vulnerable since they feature the ephemeral-static ECDH optimisation. The attack implementation facilitates recovery of the application instance’s ECDH private key. Once carried out, it allows the attacker to decrypt all previous SSL sessions from the application instance and to impersonate the server until the application restarts.

Algebraic and Algorithmic Countermeasures. Coron outlines three methods to thwart DPA attacks [11, Section 5]. In general, they seek to counteract the deterministic nature of double-and-add style scalar multiplication routines, and are therefore quite effective against the bug attack presented above.

Scalar blinding, i.e., $[k]P = [k + rn]P$ for (small) random value r , effectively randomises the execution path of the scalar multiplication algorithm. This is not “free” however: the performance overhead (and security) is proportional to the size of the random multiplier. Point blinding, i.e., $[k]P = [k](P + R) - S$, with randomly chosen $R \in E$ and $S = [k]R$ (updating both $R = [r]R$ and $S = [r]S$ for small, random r periodically), is equally effective. However, this also entails some performance overhead and is slightly more intrusive to integrate. Lastly, coordinate blinding, i.e., multiplying the projective coordinates of the accumulator by a random element of the finite field in such a way that preserves the equivalence class, effectively randomises the states of the scalar multiplication algorithm. In this case, said blinding would only need to occur at the beginning of the algorithm and hence does not entail anywhere near as significant a performance overhead. Our implementation as a patch to the `OpenSSL` source code is available from the `openssl-dev` mailing list¹².

Algorithmic countermeasures seem ineffective against the attack if they are deterministic. The Montgomery ladder [27, Section 10.3.1], for example, can resist many forms of side-channel attack due to the regular nature of operations performed; it cannot resist a variant of the attack in Section 3 however, since one can still select distinguisher points that target the control-flow and hence (iteratively) recover the scalar. For completeness, we reproduce the approach

¹² <http://marc.info/?l=openssl-dev&m=131194808413635>

Input: The point $P \in E$ and scalar $k = (k_{l-1}, \dots, k_1, k_0)$ expressed in base-2.

Output: The point $Q = [k]P \in E$.

```
 $R_0 \leftarrow \mathcal{O}$ 
 $R_1 \leftarrow P$ 
for  $i = l - 1$  downto 0 step  $-1$  do
  if  $k_i = 0$  then
     $R_1 \leftarrow R_0 + R_1$ 
     $R_0 \leftarrow [2]R_0$ 
  end
  else
     $R_0 \leftarrow R_0 + R_1$ 
     $R_1 \leftarrow [2]R_1$ 
  end
end
return  $R_0$ 
```

Algorithm 1: The Montgomery ladder algorithm for scalar multiplication.

in Algorithm 1. Notice that for every i -th iteration of the loop there are two control-flow cases possible, namely

1. if $k_i = 0$ then the assignments $R_1 \leftarrow R_0 + R_1$ and $R_0 \leftarrow [2]R_0$ occur, otherwise
2. if $k_i = 1$ then the assignments $R_0 \leftarrow R_0 + R_1$ and $R_1 \leftarrow [2]R_1$ occur.

To recover k_i , one simply selects an input P such that in the i -th iteration computation of $[2]R_0$ induces the bug yet $R_0 + R_1$ and $[2]R_1$ do not. Then, $k_i = 1$ (resp. $k_i = 0$) if the result is valid (resp. invalid) meaning after the loop terminates $R_0 \in E$ (resp. $R_0 \notin E$). Indeed, similar variants seem valid for a large class of square-and-multiply (or double-and-add) algorithms using any exponent (or scalar) representation that is not randomised (e.g., binary, m -ary, sliding window, NAF).

4 Approaches to Formal Verification

In this section we investigate whether it is realistic to use current formal verification technology to prevent similar bug attacks in open-source projects such as **OpenSSL**. We focus our analysis in two complementary aspects of this problem: first the design of efficient algorithms for carrying out the necessary numeric computations, and second checking that these algorithms are correctly implemented in machine code. The concrete arithmetic bug that we explore in this paper serves as a perfect illustration of why these two aspects should be considered separately.

A high-level specification of the procedure for modular reduction that was found to be incorrect is described in Section 2.4. Producing a concrete implementation of this procedure implies a need to refine it into a lower-level specification; the particular refinement we are analysing can be described as follows:

1. Pre-compute a table T , where the i -th element $T[i] = i \cdot p$ (for small i).
2. To reduce the integer product $z = x \cdot y$ modulo p , first construct the intermediate values S_0, S_1, \dots, S_8 based on the representation of z .
3. Write the integer sum $S = S_0 + 2S_1 + 2S_2 + S_3 + S_4 - S_5 - S_6 - S_7 - S_8$ as $S = t + 2^{256} \cdot c$.
4. Return the result $r' = t - \text{sign}(c) \cdot T[|c|] \pmod{2^{256}}$.

This highlights a subtle point: rather than a programming error, the bug is more accurately characterised as a design error. That is, the *incorrectly designed* refinement is *correctly implemented* in OpenSSL.

In the next subsections we first discuss how one can formally verify the design of such refinements using existing technology, and then viability of fully verifying implementation correctness.

4.1 Verifying Correct Refinement

The first question we consider is whether it is feasible to verify that a particular refinement is correct wrt. the associated high-level specification. In order to illustrate the techniques that can be employed at this level, we present two examples inspired in the bug described in the previous sections; each represents a refinement (invalid and valid respectively) along the lines above.

We have used the CAO domain specific language for cryptography [6] and the CAOVerif deductive verification tool [34]. The CAO language is syntactically close to C, but is equipped with type system (including, in particular, multi-precision integers) that make it straightforward to express mathematically-rich algorithms. The CAOVerif tool¹³ takes an annotated version of the program one wishes to prove correct as input (the specification of correctness is included in the annotations), and generates a set of proof obligations that need to be validated as output. If one is able to validate all the proof obligations, this implies the program meets the specification. The proof obligations can be discharged by automatic provers such as Simplify [12] or Alt-Ergo [10] or, if these fail, one can use an interactive theorem prover such as Coq [33]. The verification condition generation procedure is based on Hoare logic [18], and uses the Jessie/Frama-C [13] and Why [14] platforms as a back-end.

Failed Proof for an Incorrect Refinement. Proving that the refinement used by OpenSSL is functionally equivalent to the original specification essentially reduces to proving that steps 1, 3 and 4 compute the same result as Equation 1. To illustrate how the proof proceeds we implemented these steps in CAO, annotating the result using the CAO Specification Language [5] (closely inspired by the ANSI C Specification Language) to indicate the proof goal. The most relevant fragment is presented below, where `Prime` and `Power2` are global variables holding the values of p and 2^{256} , respectively:

¹³ A distribution of the CAOVerif tool and source code for the examples in this paper are available from <http://crypto.di.uminho.pt/CACE/>.

```

1 typedef modPower2 := mod[2**256];
2
3 /*@ requires ((0 <= sum) && (sum <= 7*(Power2-1))
4     && (Prime == 2**256 - 2**224 + 2**192 + 2**96 - 1)
5     && (Power2 == 2**256))
6     ensures ((0 <= result) && (result < Prime)
7         && (exists d : int; result + d*Prime == sum)) */
8 def modPrime(sum : int) : int {
9     def c, res : int;
10    c := sum / Power2;
11    /*@ assert c >= 0 && c <= 6 */
12    /*@ assert 0 <= (sum - c*Prime) */
13    /*@ assert (sum - c*Prime) < Power2 */
14    res := (int)((modPower2)sum - (modPower2)(c*Prime));
15    if (res >= Prime) {
16        res := res - Prime;
17    }
18    return res;
19 }

```

Ignoring the embedded annotations for the moment, `modPrime` takes the summation `sum` as input (which for simplicity and wlog. we assume to be positive), and computes a (possibly incorrect) output of `sum` modulo `Prime`. This computation is performed in three steps that mimic the `OpenSSL` implementation: 1) it calculates the (computationally inexpensive) division by `Power2`, 2) it uses the result `c` to subtract a multiple of `Prime` from the input (this operation is carried out efficiently modulo `Power2` by casting the values to an appropriate data type), and 3) the result is placed in the correct range by applying a conditional subtraction.

The annotations in the code can now be described. The specification of `modPrime` is a contract including a precondition `requires` and a post-condition `ensures`. It states that provided `sum` is in the correct range, i.e., $0 \leq \text{sum} \leq 7 \cdot (\text{Power2} - 1)$, and that the output meets the mathematical definition of the least residue of `sum` modulo `Prime`, i.e., $(0 \leq \text{res} < \text{Prime}) \wedge (\exists d \text{ st. } \text{res} + d \cdot \text{Prime} = \text{sum})$. Inside the function, a series of assertions guide the proof tool toward establishing intermediate results towards an attempted proof. For example, one is able to establish the correct range of `c` after the division. The proof fails, however, when one tries to establish that performing the subsequent calculation modulo `Power2` will produce a result that is still congruent with `sum` modulo `Prime`. In particular, one will not be able to prove (if that was the initial intuition) that $\text{sum} - c \cdot \text{Prime} < \text{Power2}$ which would be sufficient to ensure that the calculations *could* be performed modulo `Power2`.

Robust Proof for a Correct Refinement. Consider the following alternative refinement to that presented above.

```

1 /*@ requires ((0 <= sum) && (sum <= 7*(Power2-1))
2     && (Prime == 2**256 - 2**224 + 2**192 + 2**96 - 1)
3     && (Power2 == 2**256))
4     ensures ((0 <= result) && (result < Prime)
5         && (exists d : int; result + d*Prime == sum)) */
6 def modPrime(sum : int) : int {
7     def res : int := sum;
8     /*@ ghost def t : int :=0; */

```

```

9      /*@ assert res + t*Prime == sum */
10     /*@ invariant (res>=0) &&& (res == sum + t*Prime) */
11     while (res>=Prime) {
12         /*@ ghost t:=t-1; */
13         res := res - Prime;
14     }
15     /*@ assert ((0 <= res) &&& (res < Prime)
16                &&& (res + (-t)*Prime == sum)) */
17     return res;
18 }

```

This implements the “natural” refinement: it simply subtracts the `Prime` from the input until the result is in the appropriate range. In order to complete the proof, one needs to include a loop invariant that keeps track of how many times `Prime` is subtracted; to achieve this, we use a “ghost” variable `t` that is only visible to the verification tool. The annotated result can be fed to the `CAOverif` tool, which will automatically check that the program indeed meets the specification (noting that this automation relies partly on the assertions included).

4.2 Verifying Correct Implementation

We now briefly discuss how the implementation of `OpenSSL` could be formally verified, as is, with respect to a high-level specification such as the refinements discussed above. There are several factors which make this a challenging and a resource- and time-consuming task. Let us admit that our goal is to obtain a formally verified C implementation that can be compiled into executable machine code using a compiler that has itself been formally verified (the `CompCert` compiler [24] could be a candidate for this task). Then, two subtasks would most likely be defined in order to modularise the proof:

1. Implement a library of functions that carry out the necessary multi-precision integer operations efficiently, and prove that their semantics are consistent with the corresponding mathematical definitions.
2. Prove that the implementation of the higher-level function that relies on the above library is itself correct with respect to its specification, considering that it may include additional optimisations which are language-specific.

The first factor that complicates matters is the lack of support from development tools. Considering only the multi-precision integer library, one would ideally utilise a formally verified “off-the-shelf” component. Some effort towards partial provision of such a component is reported by Fischer [15]. However, even *if* such a library was fully-certified, provided the necessary functionality, *and* was publicly available, its performance versus `OpenSSL` would likely make it unattractive. The reason for this is that, although current formal verification technology may have arguably reached the maturity required to handle general-purpose software, it still struggles to deal with domain-specific programming techniques such as those adopted in the development of high-speed cryptographic software [1] (e.g., resorting to bit-oriented operations and pointer casts).

Furthermore, even equipped with appropriate tools and expertise, the goal is still an ambitious one. The authors in [1] describe an effort to prove an optimised C implementation of the RC4 stream cipher (also included in `OpenSSL`).

The proposed approach consists of first producing a non-optimised implementation of the intended functionality: this so-called *reference implementation* should be written in such a way that it can be checked for correctness by inspection, e.g., it is a direct transcription of a cryptographic standard (in our example above this could be a direct transcription of the CAO implementation of a suitable refinement replacing multi-precision integer operations with appropriate calls to library functions). The proof of correctness of the optimised implementation is then constructed by establishing functional equivalence to this *reference implementation*. However, due to the complexity of such proofs, the adopted proof strategy consists of gradually modifying the reference implementation, and proving functional equivalence in each step; this is conceptually similar to so-called game hopping proofs of security [31].

The same technique could be also applied to the OpenSSL implementation of the modular reduction operation discussed in this paper, e.g., to prove correctness of the unrolled sequence of operations that compute the summation

$$S_0 + 2S_1 + 2S_2 + S_3 + S_4 - S_5 - S_6 - S_7 - S_8$$

and express it as $t + c \cdot 2^{256}$. This process is feasible using current technology (e.g. the Frama-C framework [13]), and solutions towards providing some degree of automation are proposed. However, this would still be a labour-intensive activity involving highly qualified personnel.

This discussion is related to efforts such as those reported in [20], towards the creation of formally verified ARM implementations of elliptic curve cryptographic algorithms. However, there the emphasis is on ongoing work towards the development of a customised vertical framework for the generation of implementations targeting a specific platform, which is designed from scratch to allow top-to-bottom formal verification using the HOL [19] theorem prover. The approach starts at the highest level with a formalisation of the elliptic curve model itself, and includes at the lowest level a formalisation of the target architecture: the authors envisage a full formal verification of all compilation steps. The ambition of this initiative allows it to aim for extremely high assurance levels, but this also means that the technology used (when available) will probably be out of reach for scenarios such as those considered in this paper, namely in terms of scope, efficiency and required expertise.

5 Conclusions

This paper presents a concrete attack against ECDH-based functionality supported by OpenSSL version 0.9.8g. The attack works whenever the ECDH public key is static: this may occur either explicitly as a result of the selected cipher suite, or (partly) implicitly as a result of the (non-standard) ephemeral-static optimisation supported by OpenSSL. It is worth noting that we also considered exploiting the bug to mount invalid curve attacks [4], an approach we expand on in Appendix B: while this allowed us to bypass OpenSSL point validation

routines, it did not lead to a practical attack due largely to the nature of the bug severely limiting the number of invalid curves.

The arithmetic bug has been resolved in `OpenSSL` versions 0.9.8h and later. As a result, it is tempting to conclude that the attack does not represent a serious threat. However, vulnerable versions of the library are deployed in (at least) two major Linux distributions, namely Debian (as late as 5.0 “Lenny”) and for Ubuntu (as late as 9.10 “Karmic”). Although they selectively apply patches to the default installation, the arithmetic bug persists in both. That is, although a patch resolving the bug has been available since 2008, it has yet to permeate existing installations. This represents a concrete example of the premise that patching is no panacea for similar problems.

Whether `OpenSSL` should prevent optimisations like ephemeral-static being included or invoked is perhaps a more philosophical question aligned to a bigger picture. For example, problems relating to IPsec support for encryption-only modes of the Encapsulating Security Payload (ESP) protocol seems conceptually similar; a comprehensive overview and resulting attack is given by Paterson and Yau [29]. One can conjecture that the motivation for encryption-only ESP, like ephemeral-static ECDH, is efficiency. Given that provision of a more efficient, less secure option will inevitably lead to someone using it, our work lends weight to the argument that a better approach may be to permit *only* secure, albeit less efficient options.

While (non-)support for various options in `OpenSSL` is subjective in part, the correctness of what *is* supported is less debatable. As such, detecting bugs in a more rigorous manner represents a difficult and extremely resource- and time-consuming task if undertaken over the entire implementation. `OpenSSL` clearly differs from a formal cryptographic standard, but it represents a de facto, ubiquitous and mission-critical software component in many settings. As such, we suggest that the effort required to adopt a development strategy capable of supporting formal verification is both warranted, and an increasingly important area for future work.

Acknowledgements

This work has been supported in part by EPSRC via grant EP/H001689/1 and by project SMART, funded by ENIAC Joint Undertaking (GA 120224).

References

1. J. Bacelar Almeida, M. Barbosa, J. Sousa Pinto, and B. Vieira. Deductive verification of cryptographic software. *NASA Journal of Innovations in Systems and Software Engineering*, 6(3):203–218, 2010.
2. Public key cryptography for the financial services industry: The Elliptic Curve Digital Signature Algorithm (ECDSA). American National Standards Institute, ANSI X9.62-1998, 1998.

3. Public key cryptography for the financial services industry: Key agreement and key transport using elliptic curve cryptography. American National Standards Institute, ANSI X9.63-199x, 1999.
4. A. Antipa, D.R.L. Brown, A. Menezes, R. Struik, and S.A. Vanstone. Validation of elliptic curve public keys. In *Public Key Cryptography (PKC)*, pages 211–223. Springer-Verlag LNCS 2567, 2003.
5. M. Barbosa. CACE Deliverable D5.2: formal specification language definitions and security policy extensions, 2009. Available from <http://www.cace-project.eu>.
6. M. Barbosa, A. Moss, and D. Page. Constructive and destructive use of compilers in elliptic curve cryptography. *J. Cryptology*, 22(2):259–281, 2009.
7. I. Biehl, B. Meyer, and V. Müller. Differential fault attacks on elliptic curve cryptosystems. In *Advances in Cryptology (CRYPTO)*, pages 131–146. Springer-Verlag LNCS 1880, 2000.
8. E. Biham, Y. Carmeli, and A. Shamir. Bug attacks. In *Advances in Cryptology (CRYPTO)*, pages 221–240. Springer-Verlag LNCS 5157, 2008.
9. M. Ciet and M. Joye. Elliptic curve cryptosystems in the presence of permanent and transient faults. *Designs, Codes and Cryptography*, 36(1):33–43, 2005.
10. S. Conchon, E. Contejean, and J. Kanig. Ergo : a theorem prover for polymorphic first-order logic modulo theories, 2006. <http://ergo.lri.fr/papers/ergo.ps>.
11. J.-S. Coron. Resistance against differential power analysis for elliptic curve cryptosystems. In *Cryptographic Hardware and Embedded Systems (CHES)*, pages 292–302. Springer-Verlag LNCS 1717, 1999.
12. D. Detlefs, G. Nelson, and J.B. Saxe. Simplify: a theorem prover for program checking. *J. ACM*, 52(3):365–473, 2005.
13. J.-C. Filliâtre and C. Marché. Multi-prover verification of C programs. In *International Conference on Formal Engineering Methods (ICFEM)*, pages 15–29. Springer-Verlag LNCS 3308, 2004.
14. J.-C. Filliâtre and C. Marché. The Why/Krakatoa/Caduceus platform for deductive program verification. In *Computer Aided Verification (CAV)*, pages 173–177. Springer-Verlag LNCS 4590, 2007.
15. S. Fischer. Formal verification of a big integer library. In *Workshop on Dependable Software Systems at DATE'08*, 2008.
16. D. Hankerson, A. Menezes, and S. Vanstone. *Guide to Elliptic Curve Cryptography*. Springer-Verlag, 2004.
17. C. Herley. So long, and no thanks for the externalities: The rational rejection of security advice by users. In *New Security Paradigms Workshop (NSPW)*, pages 133–144, 2009.
18. C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12:576–580, 1969.
19. HOL interactive proof assistant for higher order logic. <http://hol.sourceforge.net/>, 2011.
20. J. Hurd. Formally verified elliptic curve cryptography, 2007. Talk at Computer Laboratory, Cambridge, UK.
21. Standard specifications for public key cryptography. The Institute of Electrical and Electronics Engineers, IEEE P1363 / D13, 1999.
22. N. Koblitz. Elliptic curve cryptosystems. *Mathematics of Computation*, 48:203–209, 1987.
23. P.C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In *Advances in Cryptology (CRYPTO)*, pages 104–113. Springer-Verlag LNCS 1109, 1996.

A.2 An Example Attack

Imagine the server generates

$$k_S = 4543A4D2C9DDD5516FD16D5498034C76D4EAB340276B6BAD8FF57756F4ECA8E6_{(16)}$$

whose NAF expansion (per the OpenSSL implementation, and reading from most-significant digit downward) is

$$\langle 1, 0, 0, 0, 1, 0, 0, 0, 5, \dots \rangle.$$

The aim of an attacker is to recover each digit in this expansion. The attack progresses iteratively: at the i -th step, we have recovered part of k_S (starting at $i = 1$ where we have recovered none of it) denoted by a and want to recover the next, target digit b . Based on a set of valid digits be \mathcal{D} , we select a distinguisher point $D_{a,b}$ for each $b \in \mathcal{D}$; the aim is for $D_{a,b}$ to be such that whp.

$$[a \parallel b \parallel d]D_{a,b} \notin E(\mathbb{F}_p)$$

and

$$[a \parallel c \parallel d]D_{a,b} \in E(\mathbb{F}_p)$$

for all $c \in \mathcal{D} \setminus \{0, b\}$ and some random padding d . Put simply, when fed as input to the server, the point $D_{a,b}$ shows whether the target digit is (resp. is not) b by virtue of triggering (resp. not triggering) the bug and hence failing (resp. completing) the handshake. Properties of the NAF expansion allow the attacker to infer (some) zero digits rather than test for them explicitly.

For brevity, we focus on the first three iterations of an attack which recovers the first three non-zero digits of the expansion above; in all cases we subscript variables with the iteration number.

Recovering the 1-st Non-Zero Digit. When $i = 1$ we have recovered none of k_S , and hence $a_1 = \emptyset$. The most-significant digit is a somewhat special case, and as such we let the set of valid digits be $\mathcal{D}_1 = \{1, 3, 5, 7\}$ and try distinguisher points $D_{a_1,b}$ for each $b \in \mathcal{D}_1 \setminus \{1\}$ as follows:

- Considering $k_C = \text{C0CA83FA12719999AB9B}_{(16)}$, the point

$$\begin{aligned} Q_C &= [k_C]G \\ &= (\text{5FFFF2227DED370604B9231F95171F5DBAB6423B94DEE5B196B92F16F58ABFF5}_{(16)} , \\ &\quad \text{1962CD97F5BFC6D6B15C6FA0854845B25A1AC9018E3861B3B6806CB680A94EE0}_{(16)}) \end{aligned}$$

means $Q_C \in E(\mathbb{F}_p)$. Having accepted this point as input, the server computes

$$\begin{aligned} R_S &= [k_S]Q_C \\ &= (\text{6910F8B73E2B43C0618CDA5B7BEA4D60DB7104DF3E2E6C769684A161A7153BDA}_{(16)} , \\ &\quad \text{DF3F4A8896DDB20A4CD3AB904BC1EF960CAEEEE08DC21788F6A4D187BB754BA2A}_{(16)}) \end{aligned}$$

which whp. will trigger the bug *if* the target digit is 7; here $R_S \in E(\mathbb{F}_p)$ so the bug was not triggered.

- Considering $k_C = \text{EFC70097CFE53647E65C}_{(16)}$, the point

$$\begin{aligned} Q_C &= [k_C]G \\ &= (647DE47A20252CE8A687FBE4BA4D6CF0E76AF06B9881443D8D292B593BCF792B_{(16)} , \\ &\quad 477222E3E6DDECCF34EA91D79D173198CD2260212BAAB6824E50AB431FF1A44F_{(16)}) \end{aligned}$$

means $Q_C \in E(\mathbb{F}_p)$. Having accepted this point as input, the server computes

$$\begin{aligned} R_S &= [k_S]Q_C \\ &= (1720DAD73FF2FA0E1AB9535DA2BAA1147CC25D25E011C4ED3551DF991A3D31DB_{(16)} , \\ &\quad 51714898D52DE3AB9A655E71B5C9B5766C4CD12C5F24E5626509E67CBA27F1EA_{(16)}) \end{aligned}$$

which whp. will trigger the bug *if* the target digit is 5; here $R_S \in E(\mathbb{F}_p)$ so the bug was not triggered.

- Considering $k_C = 5D660F1C47657027DCC1_{(16)}$, the point

$$\begin{aligned} Q_C &= [k_C]G \\ &= (6CA5B01DEDF9586C8EEE1B1C15AF67D9850FF9B1679BEFBB849E734EAF44AC6D_{(16)} , \\ &\quad DAA59E772AE041BB6E9CF3813355E807199A342A81B4AFC2E0E9636A391D6FBC_{(16)}) \end{aligned}$$

means $Q_C \in E(\mathbb{F}_p)$. Having accepted this point as input, the server computes

$$\begin{aligned} R_S &= [k_S]Q_C \\ &= (33713256536920BE874326D63D43E2ECAEC3059E98CB237F37687D0893395FE7_{(16)} , \\ &\quad E89A350CF481112A6E3B995ABC01C17E6B3AF2AE410886431B8E8A5C6F371EE0_{(16)}) \end{aligned}$$

which whp. will trigger the bug *if* the target digit is 3; here $R_S \in E(\mathbb{F}_p)$ so the bug was not triggered.

Since none of these distinguisher points triggered the bug, we conclude that the target digit is 1.

Recovering the 2-nd Non-Zero Digit. We know $a_2 = \langle 1 \rangle$ based on the result from the previous iteration, and that the rules governing NAF expansion interperse non-zero digits with zero digits. As such, we skip ahead to $i = 5$ by assuming $a_5 = \langle 1, 0, 0, 0 \rangle$, and let the set of valid digits be $\mathcal{D}_5 = \{0, \pm 1, \pm 3, \pm 5, \pm 7\}$. We can now try distinguisher points $D_{a_5, b}$ for each $b \in \mathcal{D}_5 \setminus \{0\}$ as follows:

- Considering $k_C = 67CB4AA317B8C33B906B_{(16)}$, the point

$$\begin{aligned} Q_C &= [k_C]G \\ &= (3FF0E57ED5045345C12EB58B2307BC905E4E947211E1356431BACFB52FDE2E02_{(16)} , \\ &\quad F11D3A541AE84030B11291E2A747D6D2E4B696732C3884C75DFAB95E77B99DD2_{(16)}) \end{aligned}$$

means $Q_C \in E(\mathbb{F}_p)$. Having accepted this point as input, the server computes

$$\begin{aligned} R_S &= [k_S]Q_C \\ &= (8AC7D48C4D856B004D0E7F75F65E8740DB2EAD1F6850C50BEEE5CF53DEB43FA5_{(16)} , \\ &\quad D31C0099EE064F89EA05F9A352EFAE1EF10DFB7A5D358B93725DD13BD9239193_{(16)}) \end{aligned}$$

which whp. will trigger the bug *if* the target digit is 7; here $R_S \in E(\mathbb{F}_p)$ so the bug was not triggered.

- Considering $k_C = 14D61C7F6061779552A5_{(16)}$, the point

$$\begin{aligned} Q_C &= [k_C]G \\ &= \left(\begin{array}{l} 3F5E95C9DAF7F59682C61537045BFBE1FFD1BC63D79923B18D9C35A2DEBE7B4_{(16)} , \\ 4E8E855F6493425DB1C2BBE792A7F5C7B91378EAF41119B14A324E5FC6404CF5_{(16)} \end{array} \right) \end{aligned}$$

means $Q_C \in E(\mathbb{F}_p)$. Having accepted this point as input, the server computes

$$\begin{aligned} R_S &= [k_S]Q_C \\ &= \left(\begin{array}{l} BEF9ADF98E9B6C4745568564F01E7C667B93C7A7B6770D55050D94B1C9CFE21F_{(16)} , \\ 427110775B5739E4BCF83522EE07F4D795A75FB5AAC503A998B1971215073ABA_{(16)} \end{array} \right) \end{aligned}$$

which whp. will trigger the bug *if* the target digit is 5; here $R_S \in E(\mathbb{F}_p)$ so the bug was not triggered.

- Considering $k_C = 37A1A0218E5D277BDA30_{(16)}$, the point

$$\begin{aligned} Q_C &= [k_C]G \\ &= \left(\begin{array}{l} E4CB0205C169F588DDD3A1DB9B86ED74C78640A557908148D2B9E6CD194EF908_{(16)} , \\ DFFE0677FA4F7C5B16F4E7C24A1F122ED2A5EF7054A0412DCBFE3146C23800F5_{(16)} \end{array} \right) \end{aligned}$$

means $Q_C \in E(\mathbb{F}_p)$. Having accepted this point as input, the server computes

$$\begin{aligned} R_S &= [k_S]Q_C \\ &= \left(\begin{array}{l} 8BB3F8EE6790376982E881930E79EC16CEA4407E492FC76AA2D4BB64D8FFAAEE_{(16)} , \\ BBF2F62188707618ABE15BD0AAA7DCEAAD24B3CFD8DFB1D5BF231555A568AC3A_{(16)} \end{array} \right) \end{aligned}$$

which whp. will trigger the bug *if* the target digit is 3; here $R_S \in E(\mathbb{F}_p)$ so the bug was not triggered.

- Considering $k_C = 7913E02855DCE3ED5C0E_{(16)}$, the point

$$\begin{aligned} Q_C &= [k_C]G \\ &= \left(\begin{array}{l} 296D416994A4801B9A48E8C67C98E0C05DE1C0E85D4DC676F32FEACDC4998F0E_{(16)} , \\ A91F9BE06C1D50EEB0295A35CA0F130F17EA647147626318E28AEC97F0653749_{(16)} \end{array} \right) \end{aligned}$$

means $Q_C \in E(\mathbb{F}_p)$. Having accepted this point as input, the server computes

$$\begin{aligned} R_S &= [k_S]Q_C \\ &= \left(\begin{array}{l} 1D0F27241C177385B0D5025029FABD5D5D8475DA4E267DCD177B49C63605C25A_{(16)} , \\ 2609897336FC4EC60E9C4EAE5C9D275D4E54717E5CFBF1F16AE45EBB091DE6B6_{(16)} \end{array} \right) \end{aligned}$$

which whp. will trigger the bug *if* the target digit is 1; here $R_S \notin E(\mathbb{F}_p)$ so the bug was triggered.

At this point we conclude that the target digit is 1, and stop testing digits from the digit set (i.e., skip those $b \in \{-1, -3, -5, -7\}$).

Recovering the 3-rd Non-Zero Digit. We know $a_6 = \langle 1, 0, 0, 0, 1 \rangle$ based on the result from the previous iteration, and that the rules governing NAF expansion intersperse non-zero digits with zero digits. As such, we skip ahead to $i = 9$ by assuming $a_5 = \langle 1, 0, 0, 0, 1, 0, 0, 0 \rangle$, and let the set of valid digits be $\mathcal{D}_9 = \{0, \pm 1, \pm 3, \pm 5, \pm 7\}$. We can now try distinguisher points $D_{a_9, b}$ for each $b \in \mathcal{D}_9 \setminus \{0\}$ as follows:

- Considering $k_C = 7A0443D0FCDF15E61472_{(16)}$, the point

$$\begin{aligned} Q_C &= [k_C]G \\ &= (CAE39085F1BDFB7C8A0F1B242B5793323717F0B8B39A513D1751B6EBECC3785A_{(16)} , \\ &\quad CB9E8BF58AA2CCB5C06215CDA8260C9C1A656C58EA3BECD4B52C9AB72088E447_{(16)}) \end{aligned}$$

means $Q_C \in E(\mathbb{F}_p)$. Having accepted this point as input, the server computes

$$\begin{aligned} R_S &= [k_S]Q_C \\ &= (FF79FE6F475B2C487B7BB47E06E8542EFEA6EC291B9E0DB4610B1DA7E2DDAAC8_{(16)} , \\ &\quad B9BD4501BD51361182B80350ED7BEBEE63A637C8DE6BC90CC7CA8AF244D10A3D_{(16)}) \end{aligned}$$

which whp. will trigger the bug *if* the target digit is 7; here $R_S \in E(\mathbb{F}_p)$ so the bug was not triggered.

- Considering $k_C = E96F03458B226E601130_{(16)}$, the point

$$\begin{aligned} Q_C &= [k_C]G \\ &= (CFCC858FD40C13054BFFE00AFB2910DA44D9B4A293EABE01483205D01A4BCCE5_{(16)} , \\ &\quad C7297562478FFF5832939F97D1425E578D50ABCA6D2F9A8365755D7C7B6DDA92_{(16)}) \end{aligned}$$

means $Q_C \in E(\mathbb{F}_p)$. Having accepted this point as input, the server computes

$$\begin{aligned} R_S &= [k_S]Q_C \\ &= (C221357931B3B344D1F0E4A1E7144813DF75606FD2053BB01E2B951994605AA1_{(16)} , \\ &\quad 86FBB9CA123D0A3BD2BF4E45591E59BF4BC6A59105416A75F43FEA290650CAF0_{(16)}) \end{aligned}$$

which whp. will trigger the bug *if* the target digit is 5; here $R_S \notin E(\mathbb{F}_p)$ so the bug was triggered.

At this point we conclude that the target digit is 5, and stop testing digits from the digit set (i.e., skip $b \in \{\pm 1, \pm 3, -5, -7\}$).

B An (Attempted) Invalid Curve Attack

B.1 Invalid Curve Attacks

Biehl et al. [7, Section 4.1] describe an attack on implementations of ECC-based protocols that do not validate input and output points, i.e., check whether they are valid. An invalid curve attack of this type exploits the fact that most formulae for point addition and doubling do not involve the coefficient a_6 . The following definition captures this:

Definition 2. [4, Section 4] *Let E be an elliptic curve defined over the finite field \mathbb{F}_p in Weierstraß form. Then an invalid curve (relative to E) is an elliptic curve E' defined over \mathbb{F}_p whose Weierstraß equation differs from that of E only in the coefficient a_6 .*

Consider a device that given a point P as input, computes $Q = [k]P$ using a secret, embedded scalar k , and returns Q as output. Within this setting an attacker can use invalid curves to their advantage; the idea is to submit a point $P \in E'$ and provoke computation on the invalid curve E' instead of E , where the former can provide drastically less security than the latter. In the most simple example, imagine the attacker submits points $P_i \in E'_i$ such that the order of all P_i are small and pairwise coprime. The attacker then collects $Q_i = [k]P_i$ and solves for k modulo the order of each P_i , and finally recovers k using the Chinese Remainder Theorem (CRT).

B.2 “Tricking” OpenSSL wrt. P-256

Per Section 2.3, OpenSSL prevents invalid curve attacks by validating input points: it will refuse to compute $Q = [k]P$ for the attacker when given some $P \in E'$ as input. More specifically, it checks that P satisfies the curve equation

$$\begin{aligned} y_P^2 &= x_P^3 + a_4 x_P + a_6 \pmod{p} \\ &= (x_P^2 + a_4) \cdot x_P + a_6 \pmod{p} \end{aligned} \quad (2)$$

by first computing, then comparing the Left-Hand Side (LHS) and Right-Hand Side (RHS). However, by triggering the bug during this validation process OpenSSL can be “tricked” into accepting an invalid point as input. To construct such a point, the attacker proceeds as follows:

1. Select x_P such that during the computation of $t = (x_P^2 + a_4) \cdot x_P + a_6 \pmod{p}$:
 - The step $t_0 = x_P^2 \pmod{p}$ *does not* trigger the bug.
 - The step $t_1 = (t_0 + a_4) \cdot x_P \pmod{p}$ *does* trigger the bug, i.e., the correct result would be $t_1 \pm 2^{256} \pmod{p}$.
 - The incorrect result t is a quadratic residue modulo p .
2. Compute $y_P = \sqrt{t} \pmod{p}$.

Put more simply, $P = (x_P, y_P)$ passes point validation since the LHS and RHS of the curve equation are equal; however, they are *only* equal because the bug was triggered meaning P is not valid wrt. the same curve that OpenSSL thinks it is.

The restriction associated with t_0 exists because this value is also used in the point addition and doubling formula. In this case Jacobian projective coordinates are used by OpenSSL uses, roughly following the formulae in P1363 [21, A.10]. If we want to ensure that $Q = [k]P \in E'$, there should be no incorrect computation within the scalar multiplication itself, only the point validation: if the former case were true, it would result in the loss of any sensible ECDLP since P and Q would lie on different curves.

Note that since the incorrect result for t_1 can only differ from the correct result by $\pm 2^{256}$; this means P will be on one of only two possible invalid curves, namely

$$\begin{aligned} E_{+256} : y^2 &= x^3 + a_4 x + (a_6 + 2^{256}) \\ E_{-256} : y^2 &= x^3 + a_4 x + (a_6 - 2^{256}) \end{aligned}$$

$$\begin{aligned}
|E_{+256}| &= \text{FFFFFFFF00000000FFFFFFFFFFFFFFFFDA0A4439003A5730FA6F898036B17E90}_{(16)} \\
&\approx 2^4 \cdot 2^{11} \cdot 2^{31} \cdot 2^{209} \\
x_P &= \text{FF00FFF00FF0000FF00FFFFFFFFFFFFFFFF0000FFF00FFFFFFFF0000000FFFFFFFF}_{(16)} \\
y_P &= \text{A0FE412A1B302B46C527AA56A2B995AEBB987B740A3100E47245BA95BC4C6918}_{(16)} \\
|E_{-256}| &= \text{FFFFFFFF0000000100000000000000001304C2CB870EB2102DEB81758D8933A44}_{(16)} \\
&\approx 2^2 \cdot 2^{11} \cdot 2^{14} \cdot 2^{16} \cdot 2^{57} \cdot 2^{154} \\
x_P &= \text{FFF0000000FFFFFFFF00000000000000FFF0000000FFFFFFFF0000FFFFFFFF}_{(16)} \\
y_P &= \text{E84DF4E4D075F21073DEFEE92451B349FC73304D1B79088E8DA03D0DAA050FF7}_{(16)}
\end{aligned}$$

Table 1. The two possible invalid curves (and points on said curves) relative to P-256.

The orders of these curves, and example input points that cause the fault in the implicit point validation, are given in Table 1. The curves suggest that applying an invalid curve attack `OpenSSL` per Section B.1 is not trivial: neither curve is smooth enough to allow complete recovery of k . However, they also suggest that while the normal security of the ECDLP on P-256 is 128-bits, the *actual* security when computed using the faulty algorithm within `OpenSSL` is reduced to 77-bits (corresponding to the hardness of the ECDLP on E_{-256}).