

Formally-Verified ASN.1 Protocol C-language Stack

Vadim Zaliva ¹ Nika Pona ²

¹Carnegie Mellon University

²Digamma.ai

What is ASN.1?

- At Digamma.ai we are verifying a compiler for ASN.1
- The ASN.1 is a language for defining data structures and rules for serialization and de-serialization.
- Initially we focus on a subset of ASN.1 used in the X.509 standard which defines the format of public key certificates.
- We formalize Basic Encoding Rules (BER) and Distinguished Encoding Rules (DER)

ASN.1 example of an X.509-like certificate

```
1 X509 DEFINITIONS ::= BEGIN
2
3 Certificate ::= SEQUENCE {
4   tbsCertificate      TBSCertificate,
5   signatureAlgorithm  AlgorithmIdentifier,
6   signature           BIT STRING
7 }
8
9 TBSCertificate ::= SEQUENCE {
10  version      [0] INTEGER,
11  serialNumber INTEGER,
12  signature    AlgorithmIdentifier,
13  issuer       Name,
14  subject      Name,
15  subjectPublicKeyInfo SubjectPublicKeyInfo,
16 }
17
18 SubjectPublicKeyInfo ::= SEQUENCE {
19   algorithm      AlgorithmIdentifier,
20   subjectPublicKey BIT STRING
21 }
22
23 AlgorithmIdentifier ::= SEQUENCE {
24   algorithm      OBJECT IDENTIFIER
25 }
26
27 Name ::= SEQUENCE OF SET OF SEQUENCE {
28   type      OBJECT IDENTIFIER,
29   value     ANY DEFINED BY type
30 }
31
32 END
```

An ASN.1 compiler parses ASN.1 syntax definitions and produces either a source code of a specialized protocol encoder/decoder for this data type or a run-time data for a parametric encoder/decoder.

We are verifying a mature open-source ASN.1 compiler, ASN1C (<https://github.com/vlm/asn1c>). It is well-tested and widely used. We do the verification in Coq proof assistant.

What Coq does?

In Coq you can:

- define functions and predicates
- state mathematical theorems and software specifications
- interactively develop formal proofs of theorems
- machine-check these proofs by a relatively small trusted kernel based on the Calculus of Inductive Constructions
- compile certified programs to languages like OCaml, Haskell or Scheme.

Preliminary work: traditional approach

First, we tried the traditional approach on an error-prone part of ASN.1: floating-point numbers encoding/decoding (<https://github.com/digamma-ai/asn1fpcoq>). We wrote the encoders/decoders in Coq, proved their correctness and extracted to OCaml. This approach is not very practical since the generated code is not as efficient and usable as the C code.

Therefore we decided to try out a different approach: verify the C code directly.

We rely on the work previously done for the CompCert project (<http://compcert.inria.fr/>). CompCert is a verified compiler for C, written in Coq and proved to work correctly

- We parse C code into a Coq abstract syntax tree using CompCert
- Write a specification in Coq
- Prove that the generated AST behaves according to the specification, according to semantics of C defined in CompCert

Preliminary experiments

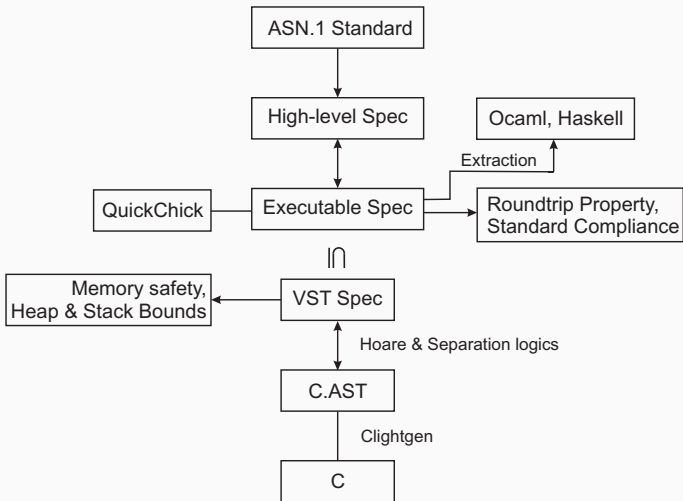
First we took a relatively simple but representative function *strtoimax* (string to integer conversion with bounds checking) from ASN1C and proved it correct using two approaches:

- proof using operational semantics defined in CompCert
- proof using separation logic defined on top of CompCert's operational semantics using Verified Software Toolchain (VST, <https://github.com/PrincetonUniversity/VST>)

During this experiment we found three bugs in this function (integer overflow, wrong memory write, semantically unintended behaviour). We saw that using VST is more practical.

Verification Architecture

We ended up with following verification architecture:



Now we explain the verification architecture on example of the boolean decoder. We focus on Basic Encoding Rules (BER).

The ASN.1 Standard says:

§8.2.1. *The contents octets shall consist of a single octet.*

§8.2.2. *If the boolean value is FALSE the octet shall be zero. If the boolean value is TRUE the octet shall have any non-zero value, as a sender's option.*

High-level spec (BOOLEAN)

- 1 **Inductive** $BER_Bool : \mathbb{B} \rightarrow list\ byte \rightarrow Prop :=$
- 2 | $False_Bool_BER : BER_Bool\ false\ [0]$
- 3 | $True_Bool_BER\ b : b \neq 0 \rightarrow BER_Bool\ true\ [b].$

BER_Bool is a relation between booleans and lists of bytes (octets) with two rules that define this relation and formalize (part of) a paragraph in the actual standard. This relation defines how a value is encoded. Then BER relation (next slide) defines how the whole packet (tag-length-value) is encoded.

High-level spec for other types

```
1 Inductive BER : asn_value → list byte → Prop :=
2 | Bool_BER b t v :
3   PrimitiveTag t → (* § 8.2.1 *)
4   BER_Bool b v →
5   BER (BOOLEAN b) (t ++[1] ++v)
6
7 | Integer_long_BER t l v z :
8   PrimitiveTag t → (* 8.3.1 *)
9   Length (length v) l → (* 10.1 *)
10  1 < length v → (* 8.3.2, case 2 *)
11  (v[0] = 255 → get_bit 0 v[1] = 0
12  ∧ v[0] = 0 → get_bit 0 v[1] = 1) → (* 8.3.2, (a) and (b) *)
13  BER_Integer z v →
14  BER (INTEGER z) (t ++l ++v)
15  ...
16
17 | Sequence_BER t l ls vs :
18  let v := flatten vs in
19  ConstructedTag t → (* 8.9.1 *)
20  Length (length v) l → (* 10.1 *)
21  (∀ n, n < length ls → BER ls[n] vs[n]) → (* 8.9.2 *)
22  BER (SEQUENCE ls) (t ++l ++v)
```

Decoder C implementation

```
1  asn_dec_rval_t
2  BOOLEAN_decode_ber(const asn_codec_ctx_t *opt_codec_ctx,
3                    const asn_TYPE_descriptor_t *td, void **bool_value,
4                    const void *buf_ptr, size_t size, int tag_mode) {
5      BOOLEAN_t *st = (BOOLEAN_t *)*bool_value;
6      asn_dec_rval_t rval;
7      ber_tlv_len_t length;
8
9      if(st == NULL) {
10         st = (BOOLEAN_t *)(*bool_value = CALLOC(1, sizeof(*st)));
11         if(st == NULL) {
12             rval.code = RC_FAIL;
13             rval.consumed = 0;
14             return rval;
15         }
16     }
17     rval = ber_check_tags(opt_codec_ctx, td, 0, buf_ptr, size,
18                         tag_mode, 0, &length, 0);
19     if(rval.code != RC_OK)
20         return rval;
21
22     buf_ptr = ((const char *)buf_ptr) + rval.consumed;
23     size -= rval.consumed;
24     if(length > (ber_tlv_len_t)size || length != 1) {
25         ASN__DECODE_FAILED;
26     }
27
28     *st = *((const uint8_t *)buf_ptr);
29
30     rval.code = RC_OK;
31     rval.consumed += length;
32
33     return rval;
34 }
```

Executable spec

Executable specification is an abstraction of the C implementation of the decoder.

```
1 Definition bool_decoder (td: TYPE_descriptor) (ls: list byte)
2   : error (byte * Z) :=
3   match ls with
4   | [] ⇒ inl FAIL
5   | _ ⇒ (consumed, expected) ← ber_check_tags td ls;
6         if Zlength ls – consumed < expected || (expected != 1)
7         then inl FAIL
8         else y ← hd (skipn consumed ls);
9         inr (y, consumed + 1)
10  end.
```

Functional correctness and the “roundtrip” property

We show that decoder is inverse of encoder.

- 1 **Theorem** *boolean_roundtrip*: $\forall td\ ls\ b\ z,$
- 2 $decoder_type\ td = BOOLEAN_t \rightarrow$
- 3 $bool_encoder\ td\ b = inr(z, ls) \rightarrow$
- 4 $bool_decoder\ td\ ls = inr(b, z).$

We prove that the executable spec encodes and decodes bytes in conformance with the high-level specification.

- 1 **Theorem** *bool_decoder_correctness*: $\forall td\ ls\ b\ z,$
- 2 $bool_decoder\ td\ ls = inr(b, z) \leftrightarrow BER(BOOLEAN\ b)(firstn\ z\ ls).$

To show C implementation correctness wrt the executable (and hence high-level spec) we prove a separation logic triple

$$P\{c\}Q$$

that given the precondition P , the execution of the C light function c terminates with the post-condition Q being true. The post-condition says that c returns the value according to the executable spec.

The memory specification uses spatial predicates $v \leftarrow p$ (“at address p there is a value v ”).

We can combine the predicates using the separating conjunction $*$: each such conjunct is true on a separate sub-heap of the memory, thus guaranteeing non-overlapping of pointers.

The precondition relates the C types such as `_asn_TYPE_descriptor_s`, `int`, `*char` of `BOOLEAN_decoder_ber` to the abstract types of Coq `TYPE_descriptor`, \mathbb{B} , `list byte` etc.

In the post-condition, we use the executable specification to state that the correct result is written in memory.

VST spec, decoder pre- and post-condition

```
1  PRE [(td : TYPE_descriptor) ← td_p *
2      (buf : list byte) ← buf_p ... *
3      bool_value_p ← bool_value_pp *
4      (res : code * Z) ← res_p *
5      if bool_value_p == null then emp else _ ← bool_value_p]
6
7  POST [ (* Unchanged memory *)
8         td ← td_p * buf ← buf_p ... *
9         (* Changed memory *)
10        EX v : val, EX ls : list val,
11         v ← bool_value_pp *
12         if v == null
13         then res ← (RC_FAIL, 0)
14         else match bool_decoder td buf with
15             | inr(r, c) ⇒ res ← (RC_OK, c) * v ← r
16             | inl FAIL ⇒ res ← (RC_FAIL, 0) * v ← ls
17         end).
```

The proof is done using so-called *forward simulation*. To prove $P\{c\}Q$:

- start assuming the precondition P
- sequentially execute statements of the function c
- each statement generates a post-condition that follows from its execution
- after executing the last statement of c , prove that the post-condition Q holds.

VST provides tactics to do most of these steps automatically. One has to provide joint postconditions for if statements and loop invariant for the loop

Lessons learned and future work

- We have the basic infrastructure in place to prove the X.509 part of ASN1C
- The memory-related parts of the proof are uniform so can be reused
- We use a layered approach to decrease the creative effort in the VST proof
- A realistic subset of C code is supported by VST
- But extensions and more automation is needed for industrial scale projects

Join us!

The project is in active development right now, but given the ambitious scope a significant effort is required for it's completion. *Digamma.ai* is committed to sponsor the initial stage of the project and we are currently looking for industry and academic partners to join us in the full ASN.1 verification endeavor.

Contact:

- Vadim Zaliva, ✉ vzaliva@cmu.edu, [🐦 @vzaliva](https://twitter.com/vzaliva)
- Nika Pona, ✉ npona@digamma.ai