

Proteus: Programmable Protocols for Censorship Circumvention

Ryan Wails^{1,2}, Rob Jansen¹, Aaron Johnson¹, and Micah Sherr²

¹U.S. Naval Research Laboratory

²Georgetown University

Abstract

We present the Proteus system for censorship circumvention. Proteus provides a programmable protocol environment in which new communication protocols can be expressed as concise and comprehensible specification files. This design allows clients and proxies to quickly respond to new censorship strategies just by installing new specification files. Proteus improves on prior programmable designs by improving host safety from malicious specifications, providing a specification language that is complete and comprehensible to non-specialists, and supporting multiple simultaneous protocols at a proxy for versioning and localization. This paper represents work in progress and provides an overview of the Proteus design, as well as examples showing that it can express existing encrypted protocols.

1 Introduction

Internet censorship is an increasingly common tool of political and social control. Consequently, anti-censorship communities have developed tools to circumvent censorship. One popular design for those tools is to relay traffic through proxies using an encrypted protocol [11, 15, 16, 21, 25, 26, 29]. However, if the censor can identify when connections are being proxied, they can block the use of those designs. Some proxy systems can be identified at the protocol level, that is, using an identifiable feature of the protocol messages, such as a header or a byte pattern [1, 24, 27, 28].

In response, Dyer et al. [6] proposed a *programmable* system for communication protocols. Their system, Marionette, provides a language and tools that make it easy to write and install new protocols at the client and proxy. This design allows new censorship methods to be quickly evaded by reconfiguring the proxy protocol. A variety of proxy protocols can be used by different proxies, making comprehensive censorship difficult to implement.

A programmable proxy system by itself does not provide a strategy to avoid detection by a censor—it only enables strategies to be quickly implemented and deployed. Censors may install blocking rules for deployed protocols, prompting the development of new protocols by evaders; this cycle is the so-called arms race interaction of censors and evaders. Evaders have some advantages in this race. As the initiators of connections, they are in a position to test and measure rules

being applied by a censor, but conversely the censor cannot easily induce a potential evader to make proxied connections. Also, the population of network users is typically large and diverse relative to the authorities and professionals designing and enforcing the censorship regime. It is typically much easier to target evasion of a relatively small set of blocking rules than it is for a censor to block a potentially large variety of circumvention strategies.

Two case studies demonstrate the usefulness of programmable protocol systems. Bock et al. [3] measured protocol filtering being applied in Iran and identify a set of rules to recognize the allowed protocols (namely, DNS, HTTP, and HTTPS). Once such rules are discovered, a programmable circumvention tool could simply distribute updated protocol specifications containing any of the allowed fingerprints. In China, measurement studies have revealed targeted blocking of Shadowsocks [2, 27], which also affects other fully encrypted protocols such as obfs4 [29] and VMess [23]. The studies reveal a protocol blacklist being applied to connections to certain destinations outside the country. The inferred rules are simple, and a programmable design would allow circumvention systems like Shadowsocks to quickly distribute protocol modifications to evade them.

Despite its potential benefits, there exist obstacles to using the Marionette system in practice. First, Marionette poses a safety risk to clients and proxies. It executes *user-specified* plugin code in a generic Python runtime environment, making its hosts vulnerable to a malicious protocol distributor that crafts the protocol files to exploit vulnerabilities or abuse privileges of the runtime. Even non-malicious protocol implementations may contain bugs that present a risk to the host machines. Accepting such a threat would give distributed proxy networks, such as the Tor network [4], a single point of failure. Second, Marionette is not expressed in a self-contained language that is both available for use today and is accessible to developers and activists. Its custom specification language is defined only implicitly by the implementation of its interpreter, and the parsing and packaging of communications data must be implemented by plugins written in a standard programming language. Third, Marionette does not support multiple protocols and version upgrades. While new protocols can be developed to respond to changes in censorship rules, clients and proxies have to synchronously upgrade to the new protocols.

To address these weaknesses, we present the *Proteus* sys-

tem. Proteus ensures safety by specifying a limited runtime system that prevents the protocol specification files from being maliciously used to exploit proxies or clients. Proteus also provides a comprehensible specification of the language for its protocol specification files. They are designed to be usable by ordinary programmers, and their message formatting component, which defines the format of individual protocol messages, requires little programming background to configure. Finally, we describe how multiple protocols can be simultaneously supported by a single Proteus proxy. As special cases of this, (1) protocol versioning can be used to respond to new censorship rules while still supporting existing clients, and (2) proxies can support clients in different locations with different strategies to evade their censors.

For a client and proxy to use Proteus to circumvent censorship, they must both be configured with the same specification files, and those files must specify a protocol that evades the techniques being applied by their censor. We do not expect the specification files to be designed by individual users. Instead, we expect that domain experts, such as the Tor Project, or activists, such as the Shadowsocks developers, will develop and distribute those files to their communities.

This paper describes work in progress on Proteus. We provide high-level descriptions of the runtime environment, a grammar for our programming language, and example protocol specifications that implement mocked versions of two existing encrypted protocols (namely, Shadowsocks and a Noise protocol [12]). Work is ongoing to fully implement Proteus and test it in target network environments. The working code repository for Proteus can be found at the following link: <https://github.com/unblockable/proteus>.

2 The Proteus System

The Proteus system is intended to enable fast reaction to a changing censorship environment. Its key design goals are (1) to enable pairwise communication, (2) to provide protocol programmability, (3) to provide safety from malicious protocol updates, and (4) to allow for graceful updates.

The basic functionality requirement is bidirectional communication between two parties. A particular focus is on enabling secure protocols that use cryptography to provide message confidentiality and integrity. While unencrypted protocols can be implemented, Proteus's library functions and parsing support are designed to facilitate cryptographic functionality, such as encryption, key exchange, and signatures.

Proteus communication protocols are programmable to allow its users to quickly adjust to changes in censorship rules and techniques. Proteus supports a wide range of different protocol state machines, message formats, and cryptographic primitives, which are commonly targets of censorship rules. Changing a protocol can easily be accomplished by updating a concise specification file which is written in a language that is designed to be familiar to programmers.

Proteus is designed to provide safety to its users by limiting the power of its execution environment thereby reducing the risk of protocol updates (relative to updating entire protocol executables). The execution environment can only interact with host operating systems through a limited set of system calls. Also, there is a limit on the memory consumed during protocol execution. Finally, the protocol specifications are expressed in a high-level language that enables inspection by the users before being installed.

2.1 Design

Proteus is designed to be used in a client-server setting. The client and proxy server communicate using a Proteus protocol designed to evade network censorship. The client is defined to be the party that initiates the connection, and the server must be running and waiting for connection attempts. Each side must possess the same Protocol Specification File (PSF) that provides the protocol specification. That PSF must be produced and distributed out-of-band, and in the setting of an adversarial censor, the PSF may need to be kept secret from the censor (for example, when specifying some distinctive but otherwise unknown protocol).

Proteus supports versioning and localization at the server. That is, the server may hold multiple PSFs and simultaneously support their multiple protocols. This feature allows the server to upgrade its protocol while remaining accessible to clients running previous protocol versions, as well as support protocols suitable for clients located in different censorship regimes. However, the method Proteus uses to choose the correct protocol requires that the supported protocols must have mutually compatible specifications to guarantee the server makes a correct protocol choice.

Multiple key setup assumptions can be used to facilitate secure communication. Keys can be provided as inputs at startup in addition to the PSFs, and then they can be used by the protocol. For example, a pre-shared symmetric key or a server public key can be provided as input by both sides to be used for encryption and authentication. Such keys must be distributed out-of-band, just as with the PSFs. Other keys may be negotiated during the protocol itself, such as ephemeral public keys or session symmetric keys, and the construction and use of those keys is specified directly in a PSF.

The system assumes that TCP is used as the underlying transport. Message delivery is assumed to be reliable and in-order. There is a notion of a connection between a pair of hosts, and it is opened by the client but may be closed by either side. The network stack may fragment messages, which should be tolerated by the protocol being used.

2.2 Abstract Model

We highlight the essential parts of the Proteus system using an *abstract model*. The Proteus abstract model consists of two




Figure 1: Relationship of the read and write buffers to a Proteus protocol. A protocol takes in data through the read buffers, and outputs data through the write buffers.

components: (1) a fixed-size execution environment Env , and (2) a protocol P to run inside of the execution environment. Protocol actions will be triggered from events defined by a set of possible events E . Each Proteus connection is handled by an independent pair of protocol instances (one instance for the client and one for the server).

The fixed-size execution environment $Env = (N, B)$ is an ordered pair determining the total state of a protocol execution: N is a positive integer that determines the size of the protocol’s global state in bytes, and B is a positive integer that determines the buffer size limit in bytes.

These parameters define the global protocol memory $G = \{0, 1, \dots, 255\}^N$ and four bounded buffers with which the protocol interacts: application read-only and write-only buffers $App^{(R)}$ and $App^{(W)}$, and network read-only and write-only buffers $Net^{(R)}$ and $Net^{(W)}$, each consisting of B bytes. The relationship of these buffers to the protocol is shown in Fig. 1.

Protocol $P = (F, \delta)$ is an ordered pair parameterized by Env . F is a finite set of functions F_1, \dots, F_k . Each function takes as input the memory and buffer state and outputs new state, i.e., $F_i : \{0, 1, \dots, 255\}^{N+4B} \rightarrow \{0, 1, \dots, 255\}^{N+4B}$. Each function is a fixed-sized boolean circuit. $\delta : E \rightarrow F$ is a dispatch function that maps each event to an event handling function.

Proteus protocols are *event driven*, which is a common programming paradigm for message passing and network protocols. Events are generated and enqueued as application and network transmissions occur. Events are processed in a loop where each event invokes an event-handling function F_i determined by δ . The event-handling loop is shown in Alg. 1.

Events are assumed to occur atomically and may be generated concurrently as the protocol is executed (e.g., an implementation of the Proteus runtime could run Alg. 1 in one thread of execution and monitor for events in another thread). The set of possible events E is given in Table 1. The most common events are EV-APP and EV-NET, which occur when new data is made available by the application or communicating party. Other events are used to handle connection initialization, termination, and errors.

2.3 Implementation

The abstract model is useful for understanding how Proteus protocols work, but does not describe how these protocols are specified or the details of the protocol runtime environ-

Algorithm 1 Main event-handling loop for Proteus protocols.

▷ Initialization:

- 1: $G \leftarrow 0^N$
- 2: $App^{(R)} \leftarrow 0^B$
- 3: $Net^{(R)} \leftarrow 0^B$
- 4: $App^{(W)} \leftarrow 0^B$
- 5: $Net^{(W)} \leftarrow 0^B$
- 6: EV-INIT is placed on the event queue.

▷ Event Processing:

- 7: **repeat**
 - 8: The next event e is popped from the event queue. Execution is paused if an event is not yet available.
 - 9: The event handler function is obtained: $f \leftarrow \delta(e)$.
 - 10: Let S be shorthand notation for the state of the execution, $S \equiv (G, App^{(R)}, Net^{(R)}, App^{(W)}, Net^{(W)})$. The event handler for e is invoked and state is updated: $S \leftarrow f(S)$.
 - 11: Data added to the application write-only buffer $App^{(W)}$ is written to the application. The buffer is reset: $App^{(W)} \leftarrow 0^B$. The same action is then applied to the network write-only buffer.
 - 12: **until** $e = \text{EV-TERM}$
-

ment. Here we describe the *Proteus language* that is used to define the set of event handling functions F_1, \dots, F_k described in the abstract model, which fully specifies a protocol. These function definitions are stored in a single source code file, the protocol specification file (PSF). In order for the language to be both simple and safe, we intentionally limited its capabilities. For example, Proteus programs have no way of dynamically managing memory. To enable complex functionalities necessary for transport protocols, such as encryption, a *standard library* of functions is provided for programs to use. Because transport protocols heavily involve message serialization and parsing, Proteus has facilities and standard library functions to simplify message formatting.

2.3.1 Proteus Language

Proteus protocols are expressed in a PSF consisting of (1) protocol message definitions (described further in § 2.3.3), (2) global state variables, and (3) event handling functions. This layout is depicted in Fig. 2. We define a custom language which is used to write Proteus protocols. The syntax of the language is designed to be familiar to Rust programmers and the language has typical low-level language semantics. A parsing expression grammar recognizing the language is given in Appendix C. The language is designed to be simple, minimal, easily edited, and interpreted at runtime. A variety of standard programming language constructs are supported, including: variable declaration and assignment; basic logical and arithmetic operations; branching execution with `if` and `match` statements; type casting; standard library function invocation; and repeated evaluation with statically-bounded `for` loops. The language is statically typed and statically

Table 1: Description of events defining the event set E .

| Event | Description |
|--------------|---|
| EV-INIT | The initialization event will always occur exactly once at the very beginning of every protocol execution. |
| EV-APP | New data was written from the application into the application read buffer App ^(R) . |
| EV-NET | New data was written from the network into the network read buffer Net ^(R) . |
| EV-TIMER | A timer expired. |
| EV-SIGQUIT | The execution process received a quit or kill signal. |
| EV-PANIC | The execution process encountered an unrecoverable error, such as an out-of-memory error. |
| EV-APP-CLOSE | The application closed its side of the connection. |
| EV-NET-CLOSE | The network closed its side of the connection. |
| EV-TERM | The final termination event. This event occurs exactly once at the very end of a protocol execution. It fires (1) immediately following EV-SIGQUIT, (2) immediately following EV-PANIC, or (3) after both the application and network connections are closed. |

Layout of a Protocol Specification File (PSF)



Figure 2: Schematic overview of PSF that consists of: (1) protocol message formats, (2) global variables used by event handlers; and (3) event handling functions.

allocated, with simple function-level lexical scoping and lifetimes (except for global variables, which have global scope and static lifetime). Listing 1 shows a simple example of code written in the Proteus language.

We intentionally limit the Proteus language to include only a small number of basis functionalities in order to promote *safety* of the language and execution environment. Specifically, we exclude: dynamic memory allocation; function or class declaration; template and macro metaprogramming¹; exceptions or exception handling; arbitrary system calls; infinite or dynamic loops; jumps; floating point arithmetic; pointer arithmetic; concurrency; and explicit memory dereferencing. Proteus programs cannot consume more than a fixed amount of memory, and procedure execution times may be measured before the procedures are invoked. Unsafe memory operations are disallowed to prevent this common source of programming errors. Many of these choices coincide with common standards for writing safety-critical code [9].

To limit host-machine access, only the narrow set of necessary system calls is allowed by the runtime. These trusted system calls related to network communication are made available through the standard library of functions available to Proteus programs, which we describe next.

¹We do allow a limited number of trusted standard library functions to be defined with template types and macros to improve code concision.

```
1 let n: u16 = 0;
2
3 for n in 1..=100 {
4   if n % 15 == 0 {
5     log("fizzbuzz");
6   } else if n % 3 == 0 {
7     log("fizz");
8   } else if n % 5 == 0 {
9     log("buzz");
10  }
11 }
```

Listing 1: A simple example showing the “fizzbuzz” program implemented in the Proteus language. The syntax closely follows that of the Rust language.

2.3.2 Standard Library

Because Proteus programs are fairly limited in what they can express, a *standard library* is defined to provide common and required functionalities for communication protocols. Standard library details and functions are further described in Appendix B. Categories of functions include:

I/O Related: These functions are used to manipulate the communication buffers. Functions include `buffer_length()`, `buffer_peek()`, `buffer_pop()`, `buffer_push()`, `buffer_close()`, and `buffer_close_all()`.

Utility: Utility functions are also provided for operations such as getting the value of an environment variable or setting a timer. Functions include `getenv()` (which retrieves the value of an environment variable), `log()`, `arm_timer()`, `disarm_timer()`, `get_timer()`, and `get_random_bytes()`.

Message Formatting: Special functions are provided to format and parse protocol messages. These functions are described further in § 2.3.3.

Cryptographic: A number of cryptographic facilities must be provided to support common operations, such as encryption and message authentication. We assume a standard set of functionalities in the standard library, such as those provided by the RustCrypto packages [20].

2.3.3 Message Formatting

Message formatting constitutes an central part of the Proteus language. The Proteus language includes *message definition* functionality, where the layout and binary encoding of protocol messages can be defined. The syntax for protocol message formats is contained in the Proteus language grammar (Appendix C). An example of a message format specification is shown in Listing 2. This example describes a protocol message called `EncryptedMessageFormat` with 3 fields: (1) `PayloadSize`, (2) `EncryptedPayload`, and (3) `MACTag`. The order of enumeration in the format specifier defines the order that these fields appear in the serialized message. Each field

```

1 DEFINE EncryptedMessageFormat
2 { NAME: PayloadSize; TYPE: u16 },
3 { NAME: EncryptedPayload;
4   TYPE: [u8; PayloadSize.value] },
5 { NAME: MACTag; TYPE: [u8; 16] };

```

Listing 2: Example protocol message definition with 3 fields—PayloadSize, EncryptedPayload, and MACTag—that are serialized by their order of declaration and types.

has a corresponding type parameter, which determines how the field is represented in binary format. *Arrays* with type t and length ℓ are denoted $[t; \ell]$. Array lengths may be concretely defined (e.g., 2 elements), or defined using a simple unambiguous expression (e.g., for the EncryptedPayload field, the size is set to PayloadSize.value, which indicates that the PayloadSize field stores the length of the field. An example of message formatting and parsing is shown below:

```

1 // Dummy values. In practice, the payload and MAC tag
2 // would be set by an encryption function.
3 let payload: [u16; 30] = [0; 30];
4 let mac: [u16; 16] = [0; 16];
5 let payload_size: u16 = 30;
6
7 // Sets the value of each field.
8 let fields: Fields = create_fields();
9 set_field(&fields, "PayloadSize", payload_size);
10 set_field(&fields, "EncryptedPayload", payload);
11 set_field(&fields, "MACTag", mac);
12
13 // The fields can then be serialized according to
14 // the EncryptedMessageFormat into a byte string.
15 let (success, b): (bool, Bytes) =
16   format(&EncryptedMessageFormat, &fields);
17
18 // This invocation parses byte string b
19 // following the EncryptedPayloadSpec format and
20 // stores the result in f2.
21 let (_, f2): (_, Fields)
22   = parse(&EncryptedMessageFormat, &b);
23
24 // Then fields can be accessed:
25 let x: u16 = 0;
26 get_field(&f2, "PayloadSize", &x);

```

Message formatting and parsing is designed to be easy and flexible. For example, in a new protocol version, a message format could be extended simply by adding new lines specifying fields’ names and types.

2.4 Versioning

Version upgrading and *localization* are important aspects of circumvention protocol design that are often overlooked. Proteus enables graceful protocol upgrades and does not require all clients and servers to update PSFs in lockstep. Instead, servers can be simultaneously provisioned with multiple protocol versions; multiple PSFs may be executed independently and in parallel using a view a single set of read buffers. State

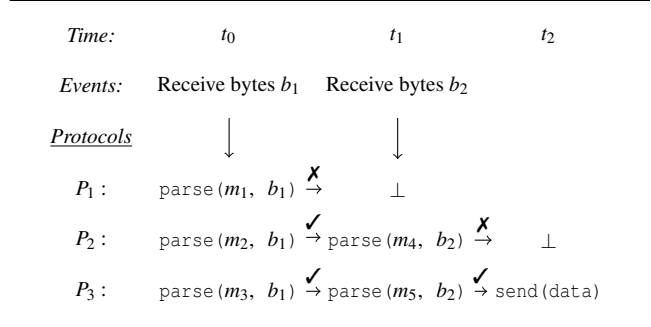


Figure 3: Diagram of protocols P_1 – P_3 simultaneously parsing incoming bytes until only one protocol P_3 remains, indicating that P_3 was the correct protocol version. \checkmark denotes when a message was successfully parsed, $\not\rightarrow$ denotes when a parse failed, and \perp denotes protocol termination.

is independently maintained for each of the running protocols. This process continues until all-but-one of the protocols have quit, or until a protocol tries to modify any one of the buffers. In the case when all-but-one have quit, the remaining protocol is determined to be the selected protocol version and continues to run. If one of the protocols modifies buffer state, then this protocol is chosen as the correct version and all other running protocol instances are immediately terminated.

This process is shown in Fig. 3. In this example, 3 protocols— P_1 , P_2 , and P_3 —are executed, each of which is configured with a separate set of message format definitions. Two events occur which correspond to receiving bytes from a client. The client (not depicted) is using protocol P_3 . Each protocol uses a different series of message formats m_i when parsing the messages. In the shown example, protocol P_1 tries parsing the first string of received bytes b_1 with an incompatible message format m_1 and quits upon failure (a parse failure can occur if, for example, a field does not contain an expected value). For protocols P_2 and P_3 , both m_2 and m_3 were compatible message formats for the first received byte string, so execution proceeds. When b_2 arrives, P_2 encounters a parsing error using format m_4 and quits, whereas P_3 ’s parsing with format m_5 is successful. At time t_2 , P_3 is the only protocol version running and is the version selected to communicate with the client.

For the Proteus versioning scheme to work as intended, Proteus protocol versions should be unambiguously determined by a client’s messages before the server is required to respond. Many transport protocols transmit a version number in the first message, which is accordant with our design.

2.5 Design Capabilities

The Proteus system contains the low-level building blocks necessary to realize high-level protocol capabilities. We now describe some of the capabilities that are commonly found in

real-world protocols and that can be achieved in Proteus.

Message Format: A protocol message is typically composed of multiple fields that contain important information to assist the receiver in parsing the message and to communicate protocol state. For example, a `length` field is often used to communicate the total length of the message. Additional information is commonly communicated in distinct message fields, such as the message type, the protocol version, a human-readable protocol greeting string, binary flags, cryptographic counters or nonces, reserved (unused) or padding bytes, message authentication codes, and application data. We are capable of expressing any number of such fields and of specifying the order in which the fields should occur within a given message by writing PSFs in the Proteus language.

Protocol Behavior: Network protocols are commonly separated into multiple protocol phases, and our language allows us to express multiple of such phases. During a *handshake* phase, specific message types are sent between the communicating parties to, for example, negotiate protocol versions, negotiate ciphersuites, and exchange cryptographic key material. The handshake phase may encompass several messages in multiple rounds of communication. Our standard library enables us to express precisely how data communicated during the handshake phase should be processed, e.g., to enable encryption. During a *data* phase, the primary focus is sending application data, possibly using an encryption method established during the handshake and possibly sending diagnostics in parallel. Finally, during a *shutdown* phase, protocols can close a connection by sending an error message or performing other termination procedures. Proteus allows us to express the logic for establishing such protocol phases.

Cryptographic Behavior: Encrypted protocols contains logic for establishing a secure communication channel. Cryptographic logic can be quite complex; for example, a ciphersuite commonly involves algorithms for key exchange, encryption, and message authentication. We support cryptographic logic through a standard library of functions, including cryptographic functions such as those supported in RustCrypto [20]. For example, Proteus allows us to express a key exchange procedure using ECDH in the Curve25519 group with the SHA256 hash function, or that encryption should be performed with a ChaCha20 stream cipher with a Poly1305 authentication tag. Functions that require auxiliary data, such as key material when constructing an ephemeral DH key, can obtain it from a peer using messages exchanged during a handshake phase as previously described.

2.6 Design Limitations

Although the Proteus system offers a large degree of flexibility due to its focus on safety and simplicity, some complex network protocols cannot be represented. For example, the file transfer protocol [14] multiplexes protocol messages over multiple connections and cannot be replicated in Pro-


teus because every client-server session is isolated to a single connection and protocol instance. Some real-world network protocols use the host’s persistent storage to maintain protocol state. TLS, for example, authenticates certificates with certificate stores located on disk. Proteus restricts system call usage from within an protocol, and hence this functionality could not be reproduced. Point-to-point transport protocols designed for censorship circumvention tend to have simple designs, leading us to believe that Proteus may be useful to program a number of protocols despite these limitations.

3 Proteus Examples

In this section, we show by example how an evader can specify and then easily modify encrypted network protocols using Proteus. We highlight salient elements of Proteus programs here and list the PSF source files in their entirety in Appendix A.

3.1 Shadowsocks

As an example, we first describe the Shadowsocks [21] obfuscation protocol as implemented in Proteus. Our implementation is *not* designed to be interoperable with Shadowsocks—it only has the same flow characteristics. To an observing third party, Shadowsocks flows have no structure and are indistinguishable from a stream of random bytes. The Shadowsocks protocol is fairly simple: each message consists of an encrypted length and an encrypted payload, where encryption is performed using an authenticated encryption with associated data (AEAD) scheme. AEAD ciphers simultaneously provide encryption and authentication, with the encryption operation outputting both a ciphertext and a *tag*, the latter of which is used by the decryption function to authenticate the ciphertext. Shadowsocks messages follow the format:



Specifying Shadowsocks in Proteus is straightforward. We first define protocol message definitions for the encrypted length (and tag) and encrypted payload (and tag). Separate message definitions are necessary since the encrypted length field needs to first be decrypted in order to determine how many bytes are required for the payload. We specify these message definitions as follows:

```
1 DEFINE EncLenFmt // encrypted length
2 { NAME: EncPayloadLen; TYPE: [u8; 2]; }
3 { NAME: EncPayloadLenTag; TYPE: [u8; 16]; };
4
5 DEFINE EncPayloadFmt // the payload
6 { NAME: EncPayload; TYPE: [u8; *]; }
7 { NAME: EncPayloadTag; TYPE: [u8; 16]; };
```

Listing 3: Message definitions for Shadowsocks

Following the Shadowsocks specification, we use two bytes for the encrypted payload length and 16 bytes for all tags. As with Shadowsocks, we use the ChaCha20 stream cipher with (16 byte) Poly1305 message authentication codes.

The PSF file also defines event handlers for the events described in Table 1:

```
1 SET_HANDLER( EV_NET, evNetRead );
2 SET_HANDLER( EV_APP, evAppRead );
3 ...
```

The main operation of our Proteus-based Shadowsocks implementation is described in the `evNetRead()` and `evAppRead()` handlers (see § A.1 for their full descriptions). `evNetRead()` computes the length of an `EncLenFmt` message, $2 + 16 = 18$ bytes, and calls `pop()` on `Net(R)` to read 18 bytes off of the network read buffer. The `parse()` function then casts those bytes into an `EncLenFmt` message:

```
1 let encLen: Fields =
2   match parse(&EncLenFmt, &encLenBytes) {
3     (true, v) => v, ... };
```

Given the resulting message, the `decrypt()` function is called to obtain the payload size, `pl` (in plaintext). The handler then reads another `pl` bytes from `Net(R)` and calls `parse()` on the returned bytes to obtain the `EncPayloadFmt` message:

```
1 let payload: Fields = match
2   parse(&EncryptedPayloadFmt, &encPayload) {
3     (true, v) => v, ... };
```

Because the length of the `EncPayload` field in the `EncPayloadFmt` message is not known before receiving and decrypting the encrypted payload length, the `*` size indicator in the message definition is necessary (see Listing 3). This tells the `parse()` function to first assign all other fields (here, just the fixed-sized `EncPayloadTag` field) before assigning the remaining bytes in the buffer to the `EncPayload` field.

Finally, the `decrypt()` function is called again to obtain the plaintext payload. The decrypted payload is then pushed to the `App(W)` buffer for reading by the application.

The `evAppRead()` event handler performs the mirror operations with respect to `evNetRead()`: it reads bytes from `App(R)` (data sent by the application) and encrypts (1) the number of bytes read, and (2) the read bytes, both using ChaCha20-Poly1305. It then calls `format()` to construct the `EncLenFmt` and `EncPayloadFmt` messages²:

```
1 let encLenSpec: [u8; *] = match format(&EncLenFmt,
2   &format!["EncPayloadLen", encLen],
3   ("EncPayloadLenTag", encLenTag) ] )
4   { (true,v) => v, ... };
5
6 let encPayloadSpec: [u8; *] =
7   match format(&EncryptedPayloadFmt,
8   &format!["EncPayload", encPayload],
9   ("EncPayloadTag", encPayloadTag) ] )
10  { (true,v) => v, ... };
```

²The `format! [...]` construct used in this example is syntactic sugar to create a `format` object with the specified fields.

The `format()` function returns the byte-representation of the messages, which are then pushed to the `Net(W)` buffer for transport over the network.

3.2 Modifying Shadowsocks

Wu et al. recently exposed a number of heuristics used by the Great Firewall (GFW) in China to detect and block Shadowsocks [27]. Essentially, the GFW looks for and blocks apparently high-entropy connections that are not TLS or HTTP. However, Wu et al. note that the GFW’s approach to blocking Shadowsocks is brittle. In particular, connections are allowed if the first 6 bytes of the first packet of a flow are all printable characters (printable bytes are in the range `0x20–0x7E`).

Modifying the Proteus implementation of Shadowsocks to bypass GFW’s censorship is thus trivial. The `EncLenFmt` message definition can be modified as follows:

```
1 DEFINE EncLenFmtV2
2 { NAME: FixedPreamble; TYPE: [u8; 6] }, // <-- New
3 { NAME: EncPayloadLen; TYPE: [u8; 2] },
4 { NAME: EncPayloadLenTag; TYPE: [u8; 16] };
```

where `FixedPreamble` will be populated with a 6 byte alphanumeric string. Additionally, the `pop()` call in `evNetRead()` needs to read 6 more bytes than in our original Shadowsocks implementation. In total, expressing the modified Shadowsocks PSF file requires only a short patch (see Listing 5 in § A.2).

Proteus makes prototyping other packet encoding strategies easy, too. If instead of printable characters, the packet’s ratio of 0s to 1s (the packet’s so-called popcount) should be altered, a biased string could be inserted into the packet’s fields. We posit that Proteus’s adaptability is well-suited for the censorship arms race. The ability to easily modify protocols’ structure enables evaders to quickly counter new changes in behavior of the censorship system.

3.3 Noise

To further illustrate the language’s versatility, we express a Noise-based [12] protocol in Proteus; see Listing 6 in § A.3. Noise is a protocol framework that provides building blocks for constructing secure cryptographic protocols. In Listing 6, we present a Proteus-based implementation of a Noise protocol in which a client with knowledge of a server’s (e.g., bridge’s) public key performs a Diffie-Hellman exchange (with server authentication) and derives an ephemeral key, which it then uses to exchange messages via an AEAD cipher. This corresponds to the `NK` handshake pattern as described in the Noise specification [12].

For brevity, we omit a full explanation of our Noise-based protocol, and instead highlight some of the core functionalities that were expressed in Proteus. As shown in Listing 6, we use built-in crypto primitives—namely, `DH()` and `HMAC()`—to implement Noise’s key chaining and derivation algorithms.

We also separate out the logic in the `evNetRead()` and `evAppRead()` handlers based on whether the protocol is in the handshake or data transmission phase. Much of the code in Listing 6 is fairly formulaic and mostly consists of sequences of calls to `parse()` and `format()`. In total, it took less than 4 hours to express a Noise-based protocol in Proteus.

4 Related Work

Programmable Obfuscation: Format-transforming encryption (FTE) is a programmable obfuscation system that takes a regular expression as input and then modifies a data stream such that it passes the regular expression [5]. A primary use-case of FTE is to create a data stream that mimics the format of well known application protocols such as HTTP. Although FTE can modify a data stream to impose a defined structure, it offers little control over protocol semantics or the statistical properties of the obfuscated traffic.

Marionette extends FTE to improve the programmability of protocol semantics and statistical traffic properties [6]. Similar to Proteus, Marionette defines protocol state machines (called *models*) which can capture the state of a channel between multiple rounds of communication and can drive responses to particular actions such as errors. Marionette uses a domain-specific language to specify a series of *templates* that will, as in FTE, insert the bytes necessary to impose a defined structure on outgoing messages. However, this language is not specified outside of the implementation of the interpreter making it difficult even for domain experts to write correct code using the language. Comparatively, the Proteus language is specified and designed to be easy to write for both domain experts and non-specialists. Furthermore, Marionette is designed such that its language calls out to plugins written in a standard programming language to implement important data processing functionality, posing significant safety risks to users and proxy operators. In contrast, the Proteus language is intentionally limited to a core set of functions necessary to implement common functionality, and this isolation improves safety and reliability of both Proteus and the protocols it runs. Finally, unlike Proteus, Marionette does not support multiple simultaneous protocols and version upgrades.

Anti-censorship researchers activists have developed other tools offering aspects of programmability [10, 22]; however, these projects tend to lack formal documentation and maturity, making a rigorous evaluation difficult.

Programmable Anonymous Communication: Flexible Anonymous Networks (FAN) is a programmable network design that separates the software architecture from deployed functionalities [18, 19]. A FAN can be programmed by compiling functionalities (e.g., adding, removing, or modifying hook functions) using LLVM into portable RISC-V object files that get packaged and distributed as a plugin and then loaded by network nodes and executed in a sandbox using

just-in-time compilation. This approach effectively changes the code that runs inside of the anonymity network nodes.

Similar to FAN, Bento is an architecture that proposes to use middleboxes to bring network function virtualization to anonymity networks [17]. Rather than modifying Tor’s internal functions as FAN does, Bento runs as a separate process, and runs arbitrary user-defined functions in a secure enclave while interacting with Tor using its control interface.

Like our approach, FAN and Bento seek to provide better modularity to more quickly adapt to new requirements. However, they both raise significant security and trust questions since a user or plugin programmer can cause arbitrary code execution on network nodes. Our approach is more isolated and measured, focusing on providing a small standard library of functions that focus on censorship circumvention protocol behavior rather than a fully general software architecture.

5 Discussion and Future Work

Proteus is compatible with several deployment models. In coordinated systems like Tor, Proteus can enable the authorities to quickly disseminate new circumvention protocols. In loosely organized systems like Shadowsocks, Proteus could foster an ecosystem of individual experimentation to evade censorship rules as they appear in different locales.

Proteus could be seamlessly swapped into several existing systems. For existing protocols that can be expressed in the Proteus language, such as obfs4 and Shadowsocks, Proteus can work in partial deployment at only the server or client side. Moreover, on the server side, it can be used to simultaneously support improved circumvention techniques and legacy clients who have not yet upgraded.

The safety of Proteus could make automatic updating desirable for systems that adopt it. Currently, in security-conscious proxy systems like Tor, updates cannot be forced on proxy operators to limit the risk of a malicious or mistaken developer. However, this limits the speed of the arms race to how fast operators can be made to install upgrades with new evasion strategies. Proteus safety features could make pushing protocol updates no more objectionable than how the Tor authorities currently push their hourly network consensuses.

Work currently ongoing in Proteus includes completing a complete specification of the language, developing a prototype implementation, and testing it in target network environments. Possible improvements to its design include the ability to create multiple TCP connections, support for UDP, and providing more support for traffic shaping through padding bytes and added delays. Another aspect of Proteus that may be improved is error handling. Allowing Proteus protocols to implement normalized or randomized responses to errors may improve its resistance to detection via active probing [8].

Availability

Proteus is actively developed at the time of this work’s publication. The Proteus source code is maintained and updated at the following link:

<https://github.com/unblockable/proteus>.

Acknowledgments

This work was partially supported by the Office of Naval Research (ONR), the Defense Advanced Research Projects Agency (DARPA) (including under Contract No. FA8750-19-C-0500), and the Georgetown University Callahan Family Professor Chair Fund. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

References

- [1] S. Afroz, D. Fifield, M. C. Tschantz, V. Paxson, and J. Tygar, “Censorship arms race: Research vs. practice,” in *HotPETs: Workshop on Hot Topics in Privacy Enhancing Technologies*, 2015. eprint: <https://petsymposium.org/2015/papers/afroz-censor-eval-hotpets2015.pdf>.
- [2] Alice, Bob, Carol, J. Beznazwy, and A. Houmansadr, “How China detects and blocks Shadowsocks,” in *ACM IMC: ACM Internet Measurement Conference*, ACM, 2020. DOI: 10.1145/3419394.3423644.
- [3] K. Bock, Y. Fax, K. Reese, J. Singh, and D. Levin, “Detecting and evading censorship-in-depth: A case study of Iran’s protocol filter,” in *FOCI: USENIX Workshop on Free and Open Communications on the Internet*, USENIX Association, 2020. eprint: <https://www.usenix.org/system/files/foci20-paper-bock.pdf>.
- [4] R. Dingledine, N. Mathewson, and P. Syverson, “Tor: The Second-Generation Onion Router,” in *USENIX Security Symposium*, USENIX Association, 2004. eprint: https://www.usenix.org/legacy/publications/library/proceedings/sec04/tech/full_papers/dingledine/dingledine.pdf.
- [5] K. P. Dyer, S. E. Coull, T. Ristenpart, and T. Shrimpton, “Protocol misidentification made easy with format-transforming encryption,” in *ACM CCS: ACM SIGSAC Conference on Computer and Communications Security*, ACM, 2013. DOI: 10.1145/2508859.2516657.
- [6] K. P. Dyer, S. E. Coull, and T. Shrimpton, “Marionette: A programmable network traffic obfuscation system,” in *USENIX Security Symposium*, USENIX Association, 2015. eprint: <https://www.usenix.org/system/files/conference/usenixsecurity15/sec15-paper-dyer.pdf>.
- [7] B. Ford, “Parsing expression grammars: A recognition-based syntactic foundation,” in *ACM POPL: ACM SIGPLAN Symposium on Principles of Programming Language*, 2004. DOI: 10.1145/964001.964011.
- [8] S. Frolov, J. Wampler, and E. Wustrow, “Detecting probe-resistant proxies,” in *NDSS: Network and Distributed Systems Security Symposium*, Internet Society, 2020. DOI: 10.14722/ndss.2020.23087.
- [9] G. J. Holzmann, “The power of 10: Rules for developing safety-critical code,” *Computer*, 2006. DOI: 10.1109/MC.2006.212.
- [10] “Operator foundation.” Accessed 2023-06-30, Operator Foundation. (2023), [Online]. Available: <https://operatorfoundation.org/>.
- [11] “Outline client.” Accessed 2023-03-15, Jigsaw. (2023), [Online]. Available: <https://github.com/Jigsaw-Code/outline-client>.
- [12] T. Perrin, *The Noise protocol framework*, version 34, 2018. eprint: <https://noiseprotocol.org/noise.pdf>.
- [13] “pest.” Accessed 2023-01-11. (2023), [Online]. Available: <https://pest.rs/>. Archived: <http://archive.today/3Gfsr>.
- [14] J. Postel and J. Reynolds, *File transfer protocol (FTP)*, Request for Comments (RFC) 959, 1985. [Online]. Available: <https://datatracker.ietf.org/doc/html/rfc959>.
- [15] “Project V.” Accessed 2023-03-15, Project V. (2023), [Online]. Available: <https://www.v2ray.com/en/>.
- [16] “Psiphon.” Accessed 2023-03-15, Psiphon Inc. (2023), [Online]. Available: <https://psiphon.ca/ur/index.html>.
- [17] M. Reininger, A. Arora, S. Herwig, N. Francino, J. Hurst, C. Garman, and D. Levin, “Bento: Safely bringing network function virtualization to Tor,” in *ACM SIGCOMM Conference*, ACM, 2021. DOI: 10.1145/3452296.3472919.
- [18] F. Rochet, O. Bonaventure, and O. Pereira, “Flexible anonymous network,” in *HotPETs: Workshop on Hot Topics in Privacy Enhancing Technologies*, 2019. eprint: <https://petsymposium.org/2019/files/hotpets/proposals/rochet-fan.pdf>.
- [19] F. Rochet and T. Elahi, “Towards flexible anonymous networks,” arXiv, 2022. DOI: 10.48550/arXiv.2203.03764.

- [20] “Rust Crypto,” Rust Crypto. (2023), [Online]. Available: <https://github.com/RustCrypto>.
- [21] “Shadowsocks.” Accessed 2023-01-11, Shadowsocks. (2023), [Online]. Available: <https://shadowsocks.org/>. Archived: <https://archive.ph/EIB3f>.
- [22] “trojan.” Accessed 2023-06-30, trojan-gfw. (2023), [Online]. Available: <https://github.com/trojan-gfw/trojan>.
- [23] “VMess.” Accessed 2023-03-15, V2Ray. (2023), [Online]. Available: <https://www.v2ray.com/en/configuration/protocols/vmess.html>.
- [24] L. Wang, K. P. Dyer, A. Akella, T. Ristenpart, and T. Shrimpton, “Seeing through network-protocol obfuscation,” in *ACM CCS: ACM SIGSAC Conference on Computer and Communications Security*, ACM, 2015. DOI: 10.1145/2810103.2813715.
- [25] B. Wiley, “Circumventing network filtering with polymorphic protocol shapeshifting,” PhD Dissertation, The University of Texas at Austin, 2016. eprint: <https://blanu.net/Dissertation.pdf>.
- [26] B. Wiley, “Dust: A blocking-resistant Internet transport protocol,” 2011. eprint: <https://www.freehaven.net/anonbib/cache/wileydust.pdf>.
- [27] M. Wu, J. Sippe, D. Sivakumar, J. Burg, P. Anderson, X. Wang, K. Bock, A. Houmansadr, D. Levin, and E. Wustrow, “How the Great Firewall of China detects and blocks fully encrypted traffic,” in *USENIX Security Symposium*, USENIX Association, 2023. eprint: <https://www.usenix.org/conference/usenixsecurity23/presentation/wu-mingshi>.
- [28] D. Xue, R. Ramesh, A. Jain, M. Kallitsis, J. A. Halderman, J. R. Crandall, and R. Ensafi, “OpenVPN is open to VPN fingerprinting,” in *USENIX Security Symposium*, USENIX Association, 2022. eprint: <https://www.usenix.org/system/files/sec22-xue-diwen.pdf>.
- [29] Yawning Angel. “Obfs4 (the obfourscator).” (Jan. 2019), [Online]. Available: <https://github.com/Yawning/obfs4/blob/master/doc/obfs4-spec.txt>.

Appendices

A Proteus Programs

This appendix contains source code listings for Proteus programs referred to in this work. Each program was syntactically checked against the Proteus grammar given in Appendix C.

A.1 Shadowsocks

```
1  /*
2  * Event handlers
3  */
4  SET_HANDLER( EV_INIT, evInit );
5  SET_HANDLER( EV_NET, evNetRead );
6  SET_HANDLER( EV_APP, evAppRead );
7  SET_HANDLER( EV_TIMER, nullHandler );
8  SET_HANDLER( *, exitHandler ); // everything else goes to exitHandler
9
10 /*
11 * Message formats
12 */
13
14 // encrypted length
15 DEFINE EncLenFmt
16 { NAME: EncPayloadLen; TYPE: [u8; 2] },
17 { NAME: EncPayloadLenTag; TYPE: [u8; 16] };
18
19 // the payload
20 DEFINE EncryptedPayloadFmt
21 { NAME: EncPayload; TYPE: [u8; *] },
22 { NAME: EncPayloadTag; TYPE: [u8; 16] };
23
24 /*
25 * Global variables
26 */
27
28 GLOBALS {
29   let mut Key :[u8;32] = [0u8; 32];
30   // initializes nonces (counters) for incoming and outgoing traffic
31   let mut OutNonce :u64 = 0u64;
32   let mut InNonce :u64 = 0u64;
33 }
34
35 /*
36 * Event handlers
37 */
38
39 fn evInit() {
40   // grab key from environment variable
41   match getenv<[u8; 32]>("chacha20_key", &global.Key) {
42     (false, _) => panic(),
43   };
44 }
45
46
47 fn evNetRead() {
48   // compute size of the EncLenFmt frame
49   let encPayloadLenSize :u64
50     = get_field_size(&EncLenFmt, "EncPayloadLen");
51
52   let encPayloadLenTagSize :u64
53     = get_field_size(&EncLenFmt, "EncPayloadLenTag");
54
55   let expectedLen :u64 = encPayloadLenSize + encPayloadLenTagSize;
56
57   let blocking: bool = true;
58
59   // block until there are at least 'expectedLen' bytes to read
60   let encLenBytes: [u8; *] = buffer_pop(&RB_net, expectedLen, blocking);
61
62   let encLen: Fields = match parse(&EncLenFmt, &encLenBytes) {
63     (true,v) => v,
64     (false,_) => {
65       buffer_close_all(); // close connection on error
66       return;
67     }
68   };
69
70   let encPayloadLen: [u8; 2] = [0u8; 2];
71   match getField<[u8; *]>(&encLen, "EncPayloadLen", &encPayloadLen) {
72     false => panic(),
73   };
74
75   let encPayloadLenTag: [u8; 16] = [0u8; 16];
76   match getField<[u8; *]>(&encLen, "EncPayloadLenTag", &encPayloadLenTag) {
77     false => panic(),
78   };
79
80   // decrypt returns (num_bytes,val) tuple
81   let (_, l) : (_, u64) = match decrypt<u64>(
82     "chacha20-poly1305",
83     &global.Key, // key
84     &encPayloadLen, // message
85     encPayloadLenSize, // message size
86     &global.InNonce, // nonce and (next line) tag
```

```
87     &encPayloadLenTag) {
88     (true,v) => v,
89     (false,_) => {
90       buffer_close_all(); // close connection on error
91       return;
92     }
93   };
94
95   global.InNonce = global.InNonce + 1u64;
96
97   // then, grab the encrypted payload and parse it
98   let encPayload: [u8; *] = buffer_pop(&RB_net, 1, blocking);
99   let payload: Fields = match parse(&EncryptedPayloadFmt, &encPayload) {
100     (true,v) => v,
101     (false,_) => panic(),
102   };
103
104   let payload_buffer: BytesMut = get_buffer(65536u64);
105   let encPayload: [u8; *] = match getField<[u8; *]>(
106     &payload_buffer, "EncPayload", &payload_buffer) {
107     false => panic(),
108   };
109
110   let tag_buffer: BytesMut = get_buffer(16u64);
111   let encPayloadTag: [u8; *] = match getField<[u8; *]>(
112     &payload_buffer, "EncPayloadTag", &tag_buffer) {
113     false => panic(),
114   };
115
116   // and decrypt it
117   let (_, plaintext) : (_, [u8; *]) = match decrypt<[u8; *]>(
118     "chacha20-poly1305",
119     &global.Key, // key
120     &encPayload,
121     &encPayloadSize,
122     &global.InNonce,
123     &encPayloadTag)
124   {
125     (true,v) => v,
126     (false,_) => {
127       buffer_close_all(); // close connection on error
128       return;
129     }
130   };
131
132   global.InNonce = global.InNonce + 1u64;
133
134   // send the results to the app
135   buffer_push(&WB_app, &plaintext);
136 }
137
138
139 fn evAppRead() {
140   // grab data from buffer (from the application)
141   let l :u16 = buffer_length(&RB_app);
142   let data :[u8; *] = buffer_pop(&RB_app, 1, false);
143
144   // encrypt the length
145   let (encLen, encLenTag): ([u8; 2], [u8; 16]) = match encrypt(
146     "chacha20-poly1305",
147     &global.Key,
148     &l,
149     2u64,
150     &global.OutNonce ) {
151     (true,ciphertext,tag) => (ciphertext, tag),
152     (false,_,_) => panic(),
153   };
154   global.OutNonce = global.OutNonce + 1u64;
155
156   // encrypt the payload
157   let (encPayload, encPayloadTag): ([u8; *], [u8; 16]) = match encrypt(
158     "chacha20-poly1305",
159     &global.Key,
160     &data,
161     1,
162     &global.OutNonce ) {
163     (true,ciphertext,tag) => (ciphertext, tag),
164     (false,_,_) => panic(),
165   };
166
167   global.OutNonce = global.OutNonce + 1u64;
168
169   // produce the frames
170   let encLenSpec: [u8; *] = match format(&EncLenFmt,
171     &format!["EncPayloadLen", encLen],
172     &format!["EncPayloadLenTag", encLenTag])
173   {
174     } {
175     (true,v) => v,
176     (false,_) => panic(),
177   };
178
179   let encPayloadSpec: [u8; *] = match format(&EncryptedPayloadFmt,
180     &format!["EncPayload", encPayload],
181     &format!["EncPayloadTag", encPayloadTag],
182     ) {
183     (true,v) => v,
184     (false,_) => panic(),
185   };
186
187   // send them on the wire
188   buffer_push(&WB_net, &encLenSpec);
189   buffer_push(&WB_net, &encPayloadSpec);
190 }
191
192 }
```

```

195 fn nullHandler() {
196 }
197
198 fn exitHandler() {
199     exit(0u32);
200 }

```

Listing 4: PSF for Shadowsocks in AEAD mode

A.2 Modifying Shadowsocks to Bypass GFW Censorship

Modifying the Proteus implementation of Shadowsocks (see Listing 4) to bypass blocking by the GFW is straightforward. Listing 5 describes the complete patch for adding a six byte alphanumeric constant ("123456") to the beginning of Shadowsocks messages.

```

1  @@ -13,6 +13,7 @@
2
3  // encrypted length
4  DEFINE EncLenFmt
5  +{ NAME: FixedPreamble; TYPE: [u8; 6] },
6  { NAME: EncPayloadLen; TYPE: [u8; 2] },
7  { NAME: EncPayloadLenTag; TYPE: [u8; 16] };
8
9  @@ -45,6 +46,8 @@
10
11
12 fn evNetRead() {
13 + let fixed_preamble_size :u64
14 +   = get_field_size(&EncLenFmt, "FixedPreamble");
15   // compute size of the EncLenFmt frame
16   let encPayloadLenSize :u64
17   = get_field_size(&EncLenFmt, "EncPayloadLen");
18  @@ -52,7 +55,8 @@
19   let encPayloadLenTagSize :u64
20   = get_field_size(&EncLenFmt, "EncPayloadLenTag");
21
22 - let expectedLen :u64 = encPayloadLenSize + encPayloadLenTagSize;
23 + let expectedLen :u64 = encPayloadLenSize +
24 +   encPayloadLenTagSize + fixed_preamble_size;
25
26   let blocking: bool = true;
27
28  @@ -169,6 +173,7 @@
29   // produce the frames
30   let encLenSpec: [u8; *] = match format (&EncLenFmt,
31   &format![
32   +   ("FixedPreamble", "123456"),
33   ("EncPayloadLen", encLen),
34   ("EncPayloadLenTag", encLenTag)
35   ]

```

Listing 5: Modifications to the Shadowsocks PSF (see Listing 4) to achieve reduced entropy

A.3 Noise

Noise [12] is a protocol framework and does not specify wire formats. We adapt Noise to a “wire” protocol by prepending a length field in front of every message.

Noise does not correspond to a particular protocol, and instead is a framework for specifying secure protocols via *handshake patterns*. We use the NK handshake pattern, which is defined as:

```

NK:
← s      (sent out-of-band)
...
→ e, es
← e, ee

```

This corresponds to the case where the client knows apriori the server’s public key (*s*) and uses it to perform a DH exchange (with server authentication) with the server.

The corresponding Proteus definition file is presented in Listing 6.

```

1  /**
2   * Event handlers
3   */
4  SET_HANDLER( EV_INIT, evInit );
5  SET_HANDLER( EV_NET, evNetRead );
6  SET_HANDLER( EV_APP, evAppRead );
7  SET_HANDLER( EV_TIMER, nullHandler );
8  SET_HANDLER( *, exitHandler ); // everything else goes to exitHandler
9
10
11 /*
12 * Message formats
13 */
14
15 // corresponds to -> e, es
16 DEFINE Handshakel
17 { NAME: InitiatorEphemeralKey; TYPE: [u8; 45] };
18
19 // corresponds to <- e, ee
20 DEFINE Handshake2
21 { NAME: EncResponderEphemeralKey; TYPE: [u8; 56] },
22 { NAME: EncResponderEphemeralKeyTag; TYPE: [u8; 16] };
23
24 // an encrypted message (after handshaking)
25 DEFINE EncryptedPayloadSpec
26 { NAME: PayloadSize; TYPE: u16 },
27 { NAME: EncPayload; TYPE: [u8; *] },
28 { NAME: EncPayloadTag; TYPE: [u8; 16] };
29
30
31 /*
32 * Global variables
33 */
34 GLOBALS {
35   let mut CompletedHandshake :bool = false;
36   let mut OutNonce :u64 = 0u64;
37   let mut InNonce :u64 = 0u64;
38   let mut IsInitiator :bool = false;
39   let mut ck :[u8;32] = [0u8; 32];
40   let mut k :[u8;32] = [0u8; 32];
41
42   // some constants we'll need later
43   let byte01 :[u8;1] = [ 0x01; 1 ];
44   let byte02 :[u8;1] = [ 0x02; 1 ];
45
46   // key material
47   let mut EphemeralDHKeyPub :[u8;56] = [0u8; 56];
48   let mut EphemeralDHKeyPri :[u8;56] = [0u8; 56];
49   let mut ServerStaticDHKeyPub :[u8;56] = [0u8; 56];
50   let mut StaticDHKeyPub :[u8;56] = [0u8; 56];
51   let mut StaticDHKeyPri :[u8;56] = [0u8; 56];
52 }
53
54 /*
55 * Event handlers
56 */
57
58 fn evInit() {
59   // let's figure out if we're the initiator or responder
60   // by looking for the is_initiator environment variable
61
62   let mut tmp: [u8;1] = [0u8; 1];
63
64   global.IsInitiator = match getenv<[u8; 1]>("is_initiator", &tmp) {
65     (true,_) => true,
66     (false,_) => false
67   };
68
69   // initialize hash according to InitializeSymmetric(...) funcn from Noise
70   let h :[u8;32] = hash(sha256, "Noise_NK_25519_ChaChaPoly_SHA256");
71   global.ck = h;
72   global.k = h;
73
74   // compute an ephemeral key
75   (global.EphemeralDHKeyPub, global.EphemeralDHKeyPri) = genDHKeyPair(56u64);
76
77   match getenv<[u8; 56]>("server_pub", &global.ServerStaticDHKeyPub) {
78     (false,_) => panic(),
79   };
80
81   if global.IsInitiator == true {
82     // compute initial key by first doing DH...
83     let input_key_material :[u8;56]
84     = match DH(global.EphemeralDHKeyPri, global.ServerStaticDHKeyPub) {
85       (true,v) => v,
86       (false,_) => panic(),
87     };
88     // update the global key according to Noise HKDF() function
89     let temp_k :[u8;32] = hmac(sha256, global.ck, input_key_material);
90     global.ck = hmac(sha256, temp_k, global.byte01);
91     global.k = hmac(sha256, concatenate(global.ck,global.byte02));
92
93   } else {
94     // we're the server, grab key pair from environment variable
95     match getenv<[u8;56]>("dhkey_pub", &global.StaticDHKeyPub) {
96       (false,_) => panic(),
97     };
98     match getenv<[u8;56]>("dhkey_pri", &global.StaticDHKeyPri) {
99       (false,_) => panic(),
100     };
101   }
102 }
103
104
105 fn evAppRead() {

```

```

106 if !global.CompletedHandshake {
107   // before we do anything else, we need to complete the handshake
108
109   if global.IsInitiator {
110
111     // send Handshakel message to responder
112     let handshake1Fields :Fields = create_fields();
113     set_field( &handshake1Fields,
114       "InitiatorEphemeralKey", global.EphemeralDHKeyPub);
115
116     let handshake1 :[u8; *] = match format( &handshake1, &handshake1Fields ) {
117       (true,v) => v,
118       (false,_) => panic(),
119     };
120
121     buffer_push( &WB_net, handshake1 );
122
123     // wait for response from responder
124     let handshake2KeySize :u64
125     = get_field_size( &handshake2, "EncResponderEphemeralKey" );
126     let handshake2TagSize :u64
127     = get_field_size( &handshake2, "EncResponderEphemeralKeyTag" );
128     let frameContents :[u8; *]
129     = match buffer_pop( &RB_net, handshake2KeySize+handshake2TagSize, true ) {
130       (true,v) => v,
131       (false,_) => panic(),
132     };
133
134     // parse response (a handshake2 message)
135     let handshake2 :Fields = match parse( &Handshake2, frameContents ) {
136       (true,v) => v,
137       (false,_) => {
138         buffer_close_all(); // close connection on error
139         return;
140       }
141     };
142
143     let encResponderEphemeralKey :BytesMut = get_buffer( 64u64 );
144
145     if !get_field( [u8; 56] > &handshake2,
146       "EncResponderEphemeralKey", &encResponderEphemeralKey ) {
147       buffer_close_all();
148       return;
149     };
150
151     let encResponderEphemeralKeyTag :BytesMut = get_buffer( 32u64 );
152
153     if !get_field( [u8; 16] > &handshake2,
154       "EncResponderEphemeralKeyTag", &encResponderEphemeralKeyTag ) {
155       buffer_close_all();
156       return;
157     };
158
159     // decrypt to get the responder's ephemeral key
160     let responderEphemeralKey :[u8; 56] = match decrypt(
161       "chacha20-poly1305",
162       global.k,
163       &encResponderEphemeralKey,
164       handshake2KeySize,
165       &global.InNonce,
166       &encResponderEphemeralKeyTag ) {
167       (true,v) => v,
168       (false,_) => {
169         buffer_close_all();
170         return;
171       }
172     };
173
174     global.InNonce = global.InNonce + 1u64;
175
176     // update the global key according to Noise HKDF() function
177     let input_key_material :[u8;56] = responderEphemeralKey;
178     let temp_k :[u8;32] = hmac(sha256, global.ck, input_key_material);
179     global.ck = hmac(sha256, temp_k, global.byte01 );
180     global.k = hmac(sha256, global.ck, global.byte02 );
181
182     global.CompletedHandshake = true;
183
184   } else {
185
186     // we're the responder, so we'll wait for the initiator to send Handshakel
187     let handshake1KeySize :u64
188     = get_field_size( &Handshake1, "InitiatorEphemeralKey" );
189
190     let frameContents : Bytes
191     = match buffer_pop( &RB_net, handshake1KeySize, true ) {
192       (true,v) => v,
193       (false,_) => panic(),
194     };
195
196     // parse response (a handshakel message)
197     let handshakel :Fields = match parse( &Handshakel, frameContents ) {
198       (true,v) => v,
199       (false,_) => {
200         buffer_close_all();
201         return;
202       }
203     };
204
205     let initiatorPK :BytesMut = get_buffer( 64u64 );
206
207     if !get_field( [u8; 56] > &Handshakel, "InitiatorEphemeralKey", &initiatorPK ) {
208       buffer_close_all();
209       return;
210     };
211
212     // compute our first DH
213     let input_key_material :[u8;56]

```

```

214     = match DH( global.StaticDHKeyPri, initiatorPK ) {
215       (true,v) => v,
216       (false,_) => panic(),
217     };
218     let temp_k :[u8;32] = hmac(sha256, global.ck, input_key_material);
219     global.ck = hmac(sha256, temp_k, global.byte01 );
220     global.k = hmac(sha256, global.ck, global.byte02 );
221
222     // send our ephemeral key, encrypted, to the initiator
223     let (encPayload, encPayloadTag) : ([u8; *], [u8; 16]) =
224     match encrypt(
225       "chacha20-poly1305",
226       &global.k,
227       &global.EphemeralDHKeyPub,
228       56u64,
229       &global.OutNonce ) {
230       (true,ciphertext,tag) => (ciphertext, tag),
231       (false,_,_) => panic(),
232     };
233
234     global.OutNonce = global.OutNonce + 1u64;
235
236     let handshake2Fields :Fields = create_fields();
237     set_field( &handshake2Fields,
238       "EncResponderEphemeralKey", encPayload );
239     set_field( &handshake2Fields,
240       "EncResponderEphemeralKeyTag", encPayloadTag );
241
242     let handshake2 :Bytes = match format(
243       &Handshake2, &handshake2Fields ) {
244       (true,v) => v,
245       (false,_) => panic(),
246     };
247
248     buffer_push( &WB_net, handshake2 );
249
250     // and compute our final global key
251     let input_key_material2 :[u8;56]
252     = match DH( global.EphemeralDHKeyPri, initiatorPK ) {
253       (true,v) => v,
254       (false,_) => panic(),
255     };
256     let temp_k :[u8;32] = hmac(sha256, global.ck, input_key_material2);
257     global.ck = hmac(sha256, temp_k, global.byte01);
258     global.k = hmac(sha256, global.ck, global.byte02);
259
260     global.CompletedHandshake = true;
261   }
262 } else {
263   // handshake completed, so send data in AEAD
264
265   // get data from app
266   let l :u64 = buffer_length( &RB_app );
267   let data :Bytes = match buffer_pop( &RB_app, l, true ) {
268     (true,v) => v,
269     (false,_) => panic(),
270   };
271
272   // encrypt the payload
273   let (encPayload, encPayloadTag) : ([u8; *], [u8; 16]) = match encrypt(
274     "chacha20-poly1305",
275     &global.k,
276     &data,
277     l,
278     &global.OutNonce ) {
279     (true,ciphertext,tag) => (ciphertext, tag),
280     (false,_,_) => panic(),
281   };
282   global.OutNonce = global.OutNonce + 1u64;
283
284   // put the ciphertext in its frame
285   let encryptedPayloadFields :Fields = create_fields();
286   set_field( &encryptedPayloadFields, "PayloadSize", l );
287   set_field( &encryptedPayloadFields, "EncPayload", encPayload );
288   set_field( &encryptedPayloadFields, "EncPayloadTag", encPayloadTag );
289   let encPayloadSpec :Bytes
290   = match format( &EncryptedPayloadSpec, &encryptedPayloadFields ) {
291     (true,v) => v,
292     (false,_) => panic(),
293   };
294
295   // and send it!
296   buffer_push( &WB_net, encPayloadSpec );
297 }
298
299 }
300
301 fn evNetRead() {
302   if !global.CompletedHandshake {
303     // before we do anything else, we need to complete the handshake
304
305     if global.IsInitiator {
306
307       // send Handshakel message to responder
308       let handshake1Fields :Fields = create_fields();
309       set_field( &handshake1Fields,
310         "InitiatorEphemeralKey", global.EphemeralDHKeyPub);
311       let handshake1 :Bytes = match format( &Handshakel, &handshake1Fields ) {
312         (true,v) => v,
313         (false,_) => panic(),
314       };
315       buffer_push( &WB_net, handshake1 );
316
317       // wait for response from responder
318       let handshake2KeySize :u64
319       = get_field_size( &handshake2, "EncResponderEphemeralKey" );
320       let handshake2TagSize :u64
321       = get_field_size( &handshake2, "EncResponderEphemeralKeyTag" );

```

```

322 let frameContents :Bytes
323 = match buffer_pop(&RB_net, handshake2KeySize+handshake2TagSize, true) {
324 (true,v) => v,
325 (false,_) => panic(),
326 };
327
328 // parse response (a handshake2 message)
329 let handshake2 :Fields = match parse(&Handshake2, frameContents) {
330 (true,v) => v,
331 (false,_) => {
332   buffer_close_all(); // close connection on error
333   return;
334 }
335 };
336
337 let encResponderEphemeralKey :BytesMut = get_buffer(64u64);
338
339 if !get_field<[u8; 56]>(&handshake2,
340 "EncResponderEphemeralKey", &encResponderEphemeralKey) {
341   buffer_close_all();
342   return;
343 };
344
345 let encResponderEphemeralKeyTag :BytesMut = get_buffer(32u64);
346
347 if !get_field<[u8; 16]>(&handshake2,
348 "EncResponderEphemeralKeyTag", &encResponderEphemeralKeyTag) {
349   buffer_close_all();
350   return;
351 };
352
353 // decrypt to get the responder's ephemeral key
354 let responderEphemeralKey : [u8; *] = match decrypt(
355 "chacha20-poly1305",
356 global.k,
357 &encResponderEphemeralKey,
358 handshake2KeySize,
359 &global.InNonce,
360 &encResponderEphemeralKeyTag) {
361 (true,v) => v,
362 (false,_) => {
363   buffer_close_all();
364   return;
365 }
366 };
367
368 global.InNonce = global.InNonce + lu64;
369
370 // update the global key according to Noise HKDF() function
371 let input_key_material : [u8; 56] = responderEphemeralKey;
372 let temp_k : [u8; 32] = hmac(sha256, global.ck, input_key_material);
373 global.ck = hmac(sha256, temp_k, global.byte01);
374 global.k = hmac(sha256, global.ck, global.byte02);
375
376 global.CompletedHandshake = true;
377
378 } else {
379
380 // we're the responder, so we'll wait for the initiator to send Handshakel
381 let handshake1KeySize :u64
382 = get_field_size(&Handshakel, "InitiatorEphemeralKey");
383
384 let frameContents :Bytes
385 = match buffer_pop(&RB_net, handshake1KeySize, true) {
386 (true,v) => v,
387 (false,_) => panic(),
388 };
389
390 // parse response (a handshakel message)
391 let handshake1 :Fields = match parse(&Handshakel, frameContents) {
392 (true,v) => v,
393 (false,_) => {
394   buffer_close_all();
395   return;
396 }
397 };
398
399 let initiatorPK :BytesMut = get_buffer(64u64);
400
401 if !get_field<[u8; 56]>(&Handshakel, "InitiatorEphemeralKey", &initiatorPK) {
402   buffer_close_all();
403   return;
404 };
405
406 // compute our first DH
407 let input_key_material : [u8; 56]
408 = match DH(global.StaticDHKeyPri, initiatorPK) {
409 (true,v) => v,
410 (false,_) => panic(),
411 };
412
413 let temp_k : [u8; 32] = hmac(sha256, global.ck, input_key_material);
414 global.ck = hmac(sha256, temp_k, global.byte01 );
415 global.k = hmac(sha256, global.ck, global.byte02 );
416
417 // send our ephemeral key, encrypted, to the initiator
418 let (encPayload, encPayloadTag) : ([u8; *], [u8; 16]) = match encrypt(
419 "chacha20-poly1305",
420 &global.k,
421 &global.EphemeralDHKeyPub,
422 56u64,
423 &global.OutNonce) {
424 (true,ciphertext,tag) => (ciphertext, tag),
425 (false,_,_) => panic(),
426 };
427
428 global.OutNonce = global.OutNonce + lu64;

```

```

429
430 let handshake2Fields :Fields = create_fields();
431 set_field( &handshake2Fields,
432 "EncResponderEphemeralKey", encPayload );
433
434 set_field( &handshake2Fields,
435 "EncResponderEphemeralKeyTag", encPayloadTag );
436
437 let handshake2 :Bytes = match format(
438 &Handshake2, &handshake2Fields ) {
439 (true,v) => v,
440 (false,_) => panic(),
441 };
442 buffer_push( &WB_net, handshake2 );
443
444 // and compute our final global key
445 let input_key_material2 : [u8; 56]
446 = match DH(global.EphemeralDHKeyPri, initiatorPK) {
447 (true,v) => v,
448 (false,_) => panic(),
449 };
450 let temp_k : [u8; 32] = hmac(sha256, global.ck, input_key_material2);
451 global.ck = hmac(sha256, temp_k, global.byte01);
452 global.k = hmac(sha256, global.ck, global.byte02);
453
454 global.CompletedHandshake = true;
455 }
456 } else {
457 // handshake completed, so grab the data off of the wire
458 let payloadSizeLen :u64
459 = get_field_size(&EncryptedPayloadSpec, "PayloadSize");
460 let encPayloadTagLen :u64
461 = get_field_size(&EncryptedPayloadSpec, "EncPayloadTag");
462
463 let l :u16 = buffer_peek(&RB_net, payloadSizeLen);
464
465 let frameContents :Bytes = match buffer_pop(
466 RB_net,
467 payloadSizeLen + 1 + encPayloadTagLen,
468 true) {
469 (true,v) => v,
470 (false,_) => panic(),
471 };
472
473 // parse it
474 let encryptedPayload : Fields = match parse(&EncryptedPayloadSpec, &frameContents) {
475 (true,v) => v,
476 (false,_) => {
477   buffer_close_all(); // close connection on error
478   return;
479 }
480 };
481 let encPayload : BytesMut = get_buffer(65535u64);
482 if !get_field<[u8; *]>(&EncryptedPayloadSpec, "EncPayload", &encPayload) {
483   panic();
484 };
485
486 let encPayloadTag : BytesMut = get_buffer(32u64);
487 if !get_field<[u8; 32]>(&EncryptedPayloadSpec, "EncPayloadTag", &encPayloadTag) {
488   panic();
489 };
490
491 // decrypt returns (num_bytes, val) tuple
492 let plaintext : [u8; *] = match decrypt(
493 "chacha20-poly1305",
494 &global.k,
495 &encPayload,
496 l,
497 &global.InNonce,
498 &encPayloadTag) {
499 (true,v) => v,
500 (false,_) => {
501   buffer_close_all(); // close connection on error
502   return;
503 }
504 };
505
506 global.InNonce = global.InNonce + lu64;
507
508 // send the results to the app
509 buffer_push( &WB_app, plaintext );
510 }
511 }
512
513 fn nullHandler() {
514 }
515
516 fn exitHandler() {
517   exit(0u32);
518 }
519 }

```

Listing 6: PSF for the NK Noise handshake pattern

B Proteus Standard Library and Runtime

This appendix contains details regarding the functions provided by the Proteus standard library and their implementation.

B.1 Functionality

```
1 SECTION - I/O RELATED
2 =====
3
4 NAME
5
6   buffer_length - get number of bytes available in a buffer
7
8 SYNOPSIS
9
10  fn buffer_length(b: &Buffer) -> usize;
11
12 DESCRIPTION
13
14  Gets the number of bytes present in a buffer.
15
16 RETURN VALUE
17
18  The number of bytes present in the buffer. No error value is specified.
19
20 -----
21
22 NAME
23
24   buffer_peek - get a copy of n bytes from a buffer without removing them
25
26 SYNOPSIS
27
28  fn buffer_peek(b: &ReadBuffer, n: usize) -> (bool, Bytes);
29
30 DESCRIPTION
31
32  Gets the first n bytes of data present in the buffer. The buffer is not
33  modified as a result of this operation.
34
35 RETURN VALUE
36
37  buffer_peek() returns a pair of values. The first element of the pair
38  indicates if the full peek could be performed (true if so, otherwise the value
39  is false). The second element contains the copied data from the buffer
40  (with length equal to the minimum of n and buffer_length(b)).
41
42 -----
43
44 NAME
45
46   buffer_pop - removes bytes from a buffer
47
48 SYNOPSIS
49
50  fn buffer_pop(b: &mut ReadBuffer, n: usize, blocking: bool) -> (bool, Bytes);
51
52 DESCRIPTION
53
54  Removes the first n bytes from buffer b. If there are fewer than n bytes
55  available, this function does not modify the buffer. If the blocking parameter
56  is set to true, this function blocks until there is n bytes of data in the
57  buffer.
58
59 RETURN VALUE
60
61  For the first return value, returns true if the pop was successful (i.e., at
62  least n bytes were removed from the buffer); false otherwise. The second
63  argument returns the bytes that were removed.
64
65 -----
66
67 NAME
68
69   buffer_push - adds bytes to a buffer.
70
71 SYNOPSIS
72
73  fn buffer_push(b: &mut WriteBuffer, data: Bytes) -> bool;
74
75 DESCRIPTION
76
77  Adds the input data to the buffer, if there is enough room in the buffer.
78  Otherwise, the buffer is not modified.
79
80 RETURN VALUE
81
82  True if the data was successfully added to the buffer; false otherwise.
83
84 -----
85
86 NAME
87
88   buffer_close - close the connection associated with a buffer.
89
90 SYNOPSIS
91
92  fn buffer_close(b: &mut Buffer);
93
94 DESCRIPTION
95
96  Closes the connection associated with the given buffer. This operation is
97  analogous to calling close on a buffer.
98
99 RETURN VALUE
100
101  N/A
102
103 -----
104
105 NAME
```

```
106
107   buffer_close_all - closes all connections
108
109 SYNOPSIS
110
111   fn buffer_close_all();
112
113 DESCRIPTION
114
115   Equivalent to calling:
116   ```
117   buffer_close(rb_app);
118   buffer_close(rb_net);
119   buffer_close(wb_app);
120   buffer_close(wb_net);
121   ```
122
123 RETURN VALUE
124
125   N/A
126
127 -----
128
129
130
131 SECTION - UTILITY FUNCTIONS
132 =====
133
134 NAME
135
136   getenv - gets an environment variable
137
138 SYNOPSIS
139
140   fn getenv<T>(name: &str, value: &mut T) -> bool;
141
142 DESCRIPTION
143
144   Gets an environment variable of type T. The value is stored in the 'value'
145   argument if the variable is defined and can be cast to type T.
146
147 RETURN VALUE
148
149   Returns true if the environment variable was successfully stored in the value
150   argument; false otherwise.
151
152 -----
153
154 NAME
155
156   get_random_bytes - generates a number of random bytes
157                       (not suitable for cryptographic use)
158
159 SYNOPSIS
160
161   fn get_random_bytes(n: usize) -> Bytes;
162
163 DESCRIPTION
164
165   Generates n uniformly random bytes and returns them. The bytes are not
166   necessarily cryptographically strong.
167
168 RETURN VALUE
169
170   The n randomly sampled bytes.
171
172 -----
173
174 NAME
175
176   log - logs a string
177
178 SYNOPSIS
179
180   fn log(line: &str);
181
182 DESCRIPTION
183
184   Writes the input line out to the system log (defined as stderr).
185
186 RETURN VALUE
187
188   N/A
189
190 -----
191
192 NAME
193
194   panic - exits the program due to an error
195
196 SYNOPSIS
197
198   fn panic();
199
200 DESCRIPTION
201
202   Closes the program and network connections associated with the program
203
204 RETURN VALUE
205
206   N/A
207
208 -----
209
210 NAME
211
212   exit - exits the program cleanly
213
```

```

214 SYNOPSIS
215
216     fn exit(exit_code: u32);
217
218 DESCRIPTION
219
220 Closes the program and network connections associated with the program,
221 returning the specified exit code.
222
223 RETURN VALUE
224
225     N/A
226
227 -----
228
229 NAME
230
231     arm_timer - adds a timer event to the queue
232
233 SYNOPSIS
234
235     fn arm_timer(k: u8, t: usize);
236
237 DESCRIPTION
238
239 Creates a timer with number k that expires t seconds from the calling time.
240 If called with a number of a timer that was already set, this function resets
241 the timer.
242
243 RETURN VALUE
244
245     N/A
246
247 -----
248
249 NAME
250
251     disarm_timer - disables a set timer
252
253 SYNOPSIS
254
255     fn disarm_timer(k: u8);
256
257 DESCRIPTION
258
259 Disables the kth timer, if it was previously armed. Otherwise, this function
260 has no effect.
261
262 RETURN VALUE
263
264     N/A
265
266 -----
267
268 NAME
269
270     get_timer - get the number of the last timer that expired
271
272 SYNOPSIS
273
274     fn get_timer() -> u8;
275
276 DESCRIPTION
277
278 On a timer expired event, this function can be used to get the number of the
279 timer that expired.
280
281 RETURN VALUE
282
283     The number of the last timer that expired.
284
285 -----
286
287 NAME
288
289     concatenate - join two byte objects together
290
291 SYNOPSIS
292
293     fn concatenate(b1: &Bytes, b2: &Bytes) -> Bytes;
294
295 DESCRIPTION
296
297 The output is a copy of b1 and b2 joined together in sequence.
298
299 RETURN VALUE
300
301     b1 || b2
302
303 -----
304
305 NAME
306
307     get_buffer - gets a new mutable buffer
308
309 SYNOPSIS
310
311     fn get_buffer(capacity: usize) -> BytesMut;
312
313 DESCRIPTION
314
315 Returns an empty BytesMut container with the specifier capacity.
316
317 RETURN VALUE
318
319     An empty BytesMut container.
320
321 -----

```

```

322
323
324 SECTION - PROTOCOL MESSAGE MANIPULATION
325 =====
326
327 NAME
328
329     create_fields - create a new empty field object
330
331 SYNOPSIS
332
333     create_fields() -> Fields;
334
335 DESCRIPTION
336
337 Creates an initialized, empty Fields object. The Fields object is used to
338 store and retrieve message field values by name for message formatting.
339
340 RETURN VALUE
341
342     A new Fields object.
343
344 -----
345
346 NAME
347
348     set_field - set a field value
349
350 SYNOPSIS
351
352     set_field<T>(fields: &mut Fields, name: &str, value: T);
353
354 DESCRIPTION
355
356 Sets the field with the given name to be of type T and the given value.
357
358 RETURN VALUE
359
360     N/A
361
362 -----
363
364 NAME
365
366     get_field - get a field value
367
368 SYNOPSIS
369
370     get_field<T>(fields: &Fields, name: &str, value: &mut T) -> bool;
371
372 DESCRIPTION
373
374 Gets the value of the field with the specified name and type and stores it in
375 the value argument. Fails on type mismatch or if the name was not set.
376
377 RETURN VALUE
378
379     true if the value fetch was successful; false otherwise.
380
381 -----
382
383 NAME
384
385     get_field_size - gets the size of a field in a message format.
386
387 SYNOPSIS
388
389     get_field_size<T>(format: &MsgFormat, name: &str) -> usize;
390
391 DESCRIPTION
392
393 Returns the statically-defined size of a field with the given name.
394
395 RETURN VALUE
396
397     The size of the field, or 0 if the field was not present or defined to have
398     variable size.
399
400 -----
401
402 NAME
403
404     format - try to create a formatted byte string
405
406 SYNOPSIS
407
408     format(format: &MsgFormat, fields: &Fields) -> (bool, Bytes);
409
410 DESCRIPTION
411
412 Attempts to format a byte string according to the specified message format and
413 included field values.
414
415 RETURN VALUE
416
417     (true, b) for a formatted byte string b if formatting was successful.
418     (false, empty bytes) if formatting was not successful.
419
420 -----
421
422 NAME
423
424     parse - try to parse a byte string into the specified fields.
425
426 SYNOPSIS
427
428     parse(format: &MsgFormat, data: &bytes) -> (bool, Fields);
429

```



```

430 DESCRIPTION
431
432 Inverse function of format. Organizes the byte string into fields that can
433 be retrieved with the get_field function.
434
435 RETURN VALUE
436
437 For the first return value, true if the message could be successfully parsed;
438 false otherwise. If the parse was successful, the returned Fields object will
439 contain the parsed fields. Otherwise, it will be empty.
440
441 -----
442
443 SECTION - CRYPTOGRAPHIC FUNCTIONS
444 -----
445
446 We assume that ProtoSpec supports a standard set of cryptographic
447 functionalities, for example, those specified in the RustCrypto library
448 <https://github.com/RustCrypto>.

```

Listing 7: Listing of standard library functions

B.2 Implementation Details

Here we give brief, disparate remarks on implementation details related to the standard library and Proteus runtime.

Proteus programs must run using a fixed amount of memory for execution safety; however, some of the standard library functions have seemingly dynamic behavior. The standard library implementation must either (1) statically allocate all needed memory at initialization, or (2) monitor used memory, reallocating when needed but never exceeding a threshold.

Some standard library functions have the capability to *block* execution. Specifically, the network I/O function `buffer_pop()` exposes a parameter that causes execution to block until data is received. Blocking behavior is somewhat at odds with Proteus’s event-driven model; we assume that in the case of `buffer_pop()` that the blocking call “intercepts” incoming EV-NET events out-of-order. More general issues still exist, for example, if a connection is closed during a blocking call, which may lead to a deadlock. We are still exploring ways to achieve a balance between different network programming paradigms that leads to easy programming.

C Proteus Grammar

The parsing expression grammar (PEG) [7] recognizing the Proteus language is given in Listing 8. The grammar is written for the `pest` library [13], which is a Rust package used for implementing performant parsers from PEGs.

```

1  WHITESPACE = _ { " " | "\t" | NEWLINE }
2  COMMENT = _ { "//" - (!"n" - ANY)* | BLOCK_COMMENT }
3  BLOCK_COMMENT = _ { /* - (!/* - ANY)* - */ }
4
5  identifier = @{ ("_"|ASCII_ALPHA)~("_"|ASCII_ALPHANUMERIC)* }
6  dotted_identifier = { identifier ~ "." ~ identifier* }
7  compound_identifier = { "(" ~ dotted_identifier ~ ("," ~ dotted_identifier)+ ~ ")" }
8  basic_or_compound_identifier = { dotted_identifier | compound_identifier }
9
10 numeric_type = { "u8" | "u16" | "u32" | "u64" | "i8" | "i16" | "i32" | "i64" }
11 basic_type = { numeric_type | "bool" | "char" }
12 concrete_array_type = { "[" ~ type ~ "]" ~ numeric_literal ~ "]" }
13 dynamic_array_type = { "[" ~ type ~ "]" ~ "*" ~ "]" }
14 flexible_array_type = { "[" ~ type ~ "]" ~ "*" ~ disj ~ "]" }
15 array_type = { concrete_array_type | dynamic_array_type }
16 custom_type = { identifier }
17 type = { basic_type | array_type | custom_type }
18 compound_type = { "(" ~ type ~ ("," ~ type)+ ~ ")" }
19 basic_or_compound_type = { type | compound_type }
20
21 template_type = { "<" ~ basic_or_compound_type ~ ("," ~ basic_or_compound_type)* ~ ">" }
22
23 numeric_literal = @{ ASCII_DIGIT+ }
24 hex_literal = @{ "0x" - ASCII_HEX_DIGIT+ }
25 typed_numeric_literal = @{ hex_literal | (numeric_literal-numeric_type) }
26
27 string_literal = @{ "\"" - inner - "\"" }
28 inner = @{ char* }
29 char = {

```

```

30 !("\\" | "\\") ~ ANY
31 | "\\\" - ("\" | "\\\" | "/" | "b" | "f" | "n" | "r" | "t")
32 | "\\\" - ("u" - ASCII_HEX_DIGIT{4})
33 }
34
35 basic_literal = { "true" | "false" | typed_numeric_literal | string_literal
36 | numeric_literal }
37
38 array_literal = { "[" ~ literal ~ ";" ~ numeric_literal ~ "]" }
39 literal = { basic_literal | array_literal | compound_literal }
40 compound_literal = { "(" ~ literal ~ ("," ~ literal)+ ~ ")" }
41 basic_or_compound_literal = { basic_literal | compound_literal }
42
43 function_literal = { identifier ~ template_type? ~ (( "(" ~ disj? ~ ")" ) | compound_disj ) }
44
45 loc = { (dotted_identifier ~ "[" ~ disj ~ "]" ) | basic_or_compound_identifier }
46
47 disj = { (conj ~ ("|" ~ conj)* ) }
48 compound_disj = { "(" ~ disj ~ ("," ~ disj)+ ~ ")" }
49 basic_or_compound_disj = { disj | compound_disj }
50 conj = { bit_or ~ ("&&" ~ bit_or)* }
51 bit_or = { bit_xor ~ ("|" ~ bit_xor)* }
52 bit_xor = { bit_and ~ ("^" ~ bit_and)* }
53 bit_and = { equal ~ ("&" ~ equal)* }
54 equal = { rel ~ ("==" | "!=") ~ rel* }
55 rel = { bit_shift ~ ( "<" | ">" | ">=" | "<=" ) ~ bit_shift* }
56 bit_shift = { sum ~ ("<<" | ">>" ~ sum)* }
57 sum = { product ~ ("+" | "-" ~ product)* }
58 product = { unary ~ ("*" | "/" | "%" ~ unary)* }
59 unary = { ("&" ~ unary) | ("!" ~ unary) | ("-" ~ unary) | factor }
60 factor = { function_literal | format_comprehension | "(" ~ disj ~ ")" | loc | literal }
61
62 block = { "(" ~ (stmt ~ ";" )* ~ (stmt ~ ";" )? ~ ")" }
63
64 match_arm = {
65 (basic_or_compound_literal|basic_or_compound_identifier) ~
66 ">" ~ (basic_or_compound_disj|block)
67 }
68
69 match_stmt = { "match" ~ basic_or_compound_disj ~
70 "(" ~ match_arm ~ ("," ~ match_arm)* ~ ")" }
71 }
72
73 if_stmt = { "if" ~ basic_or_compound_disj ~ block ~
74 ("else" ~ "if" ~ basic_or_compound_disj ~ block)* ~
75 ("else" ~ block)? }
76 }
77
78 assignment_rhs = { if_stmt | match_stmt | block | basic_or_compound_disj }
79
80 assignment_stmt = { (loc ~ "=" ~ assignment_rhs ) }
81 decl_stmt = { "let" ~ "mut"? ~ basic_or_compound_identifier ~
82 "=" ~ basic_or_compound_type ~ ("=" ~ assignment_rhs)? }
83
84 return_stmt = { "return" ~ basic_or_compound_disj? }
85
86 format_tuple = { "(" ~ string_literal ~ "," ~ basic_or_compound_disj ~ ")" }
87
88 format_comprehension = {
89 "Format!" ~ "[" ~ format_tuple ~ ("," ~ format_tuple)* ~ "]" ~
90 "]"
91 }
92
93 for_stmt = {
94 "For" ~ identifier ~ "in" ~ numeric_literal ~ ".." ~ "="? ~ numeric_literal ~
95 block
96 }
97
98 stmt = {
99 | decl_stmt
100 | assignment_stmt
101 | if_stmt
102 | match_stmt
103 | for_stmt
104 | return_stmt
105 | block
106 | basic_or_compound_disj
107 }
108
109 handler_stmt = {
110 "SET_HANDLER" ~ "(" ~ ( ("*" | identifier) ~ ("," ~ identifier ~ ")")
111 )
112 }
113
114 handler_part = { (handler_stmt ~ ";")* }
115
116 msg_format_part = { (msg_format_stmt ~ ";")* }
117 msg_field_name = { "NAME" ~ ":" ~ identifier }
118 msg_field_type = { "TYPE" ~ ":" ~ (basic_type | array_type | flexible_array_type) }
119 msg_format_stmt = { "[" ~ msg_field_name ~ "," ~ msg_field_type ~ "]" }
120
121 global_part = { ("GLOBALS" ~ "(" ~ (decl_stmt ~ ";")* ~ ")" )? }
122
123 fn_def = { "fn" ~ identifier ~ "(" ~ ("," ~ type)* ~ ")" ~ block }
124 fn_part = { fn_def* }
125
126 psf = {
127 SOI ~
128 handler_part ~
129 msg_format_part ~
130 global_part ~
131 fn_part ~
132 EOI
133 }

```

Listing 8: Parsing expression grammar recognizing the Proteus language.