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 blocklist 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 contains 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 selfcontained 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 net 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

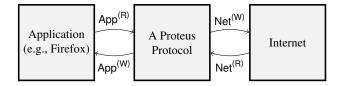


Figure 1: Relationship of the read and write buffers to a Proteus protocol. A protocol takes in data through the read buffers, and outputs data through the write buffers.

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 $G = \{0, 1, ..., 255\}^N$ 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, \ldots, F_k . Each function takes as input the memory and buffer state and outputs new state, i.e., $F_i : \{0, 1, \ldots, 255\}^{N+4B} \rightarrow \{0, 1, \ldots, 255\}^{N+4B}$. 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:

▷ Event Processing:

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) $\leftarrow 0^B$
- 5. Net $\leftarrow 0$
- 6: EV-INIT is placed on the event queue.

7: repeat

- 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: $S \leftarrow f(S)$.
- Data added to the application write-only buffer App^(W) is written to the application. The buffer is reset: App^(W) ← 0^B. The same action is then applied to the network write-only buffer.

12: until e = EV-TERM

ment. Here we describe the *Proteus language* that is used to define the set of event handling functions F_1, \ldots, 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 $\S 2.3.3$), (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 execu- tion
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 ex- actly 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)

Message Definitions		
Global State Variables		
Event Handlers		

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 metaprogrammming¹; 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.

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

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

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

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;
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);
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) =
    format(&EncryptedMessageFormat, &fields);
16
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

Time:	t_0	t_1		<i>t</i> ₂
Events:	Receive bytes b_1	Receive byte	es b_2	
<u>Protocols</u>	\downarrow	\downarrow		
P_1 :	parse (m_1, b_1)	× → ⊥		
P_2 :	parse(<i>m</i> ₂ , <i>b</i> ₁) -	∕ → parse(<i>m</i> 4,	$b_2) \xrightarrow{\mathbf{X}}$	\perp
P_3 :	parse(<i>m</i> ₃ , <i>b</i> ₁) -	\rightarrow parse(m_5 ,	$b_2) \xrightarrow{\checkmark} sen$	d(data)

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, \varkappa denotes when a parsed 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 Proteus 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:

 ← 2B →	└── 16B ──── >	K→ Var →	I6B>
Encrypted payload length	Enc. payload length auth tag	Encrypted payload	Enc. payload auth tag

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:

```
DEFINE EncLenFmt // encrypted length
{
    NAME: EncPayloadLen; TYPE: [u8; 2]; }
    {
    NAME: EncPayloadLenTag; TYPE: [u8; 16]; };

    DEFINE EncPayloadFmt // the payload
    { NAME: EncPayload; TYPE: [u8; *]; }
    { 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 <code>evNetRead()</code> and <code>evAppRead()</code> handlers (see § A.1 for their full descriptions). <code>evNetRead()</code> computes the length of an <code>EncLenFmt</code> message, 2 + 16 = 18 bytes, and calls <code>pop()</code> on Net^(R) to read 18 bytes off of the network read buffer. The <code>parse()</code> function then casts those bytes into an <code>EncLenFmt</code> 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,
       &format![("EncPayloadLen", encLen),
2
           ("EncPayloadLenTag", encLenTag) ] )
3
           { (true, v) => v, ... };
4
5
  let encPayloadSpec: [u8; *] =
6
     match format(&EncryptedPayloadFmt,
7
8
       &format![("EncPayload", encPayload),
           ("EncPayloadTag", encPayloadTag) ])
9
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] };</pre>
```

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 usecase 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

- 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://pet symposium.org/2015/papers/afroz-censor-ev al-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.us enix.org/system/files/foci20-paper-bock.p df.
- [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_paper s/dingledine/dingledine.pdf.
- [5] K. P. Dyer, S. E. Coull, T. Ristenpart, and T. Shrimpton, "Protocol misidentification made easy with formattransforming 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/f iles/conference/usenixsecurity15/sec15-pa per-dyer.pdf.
- [7] B. Ford, "Parsing expression grammars: A recognitionbased 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: 10.11 09/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/inde x.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.220 3.03764.

- [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/trojangfw/trojan.
- [23] "VMess." Accessed 2023-03-15, V2Ray. (2023), [Online]. Available: https://www.v2ray.com/en/conf iguration/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: h ttps://blanu.net/Dissertation.pdf.

- [26] B. Wiley, "Dust: A blocking-resistant Internet transport protocol," 2011. eprint: https://www.freehav en.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: http s://www.usenix.org/conference/usenixsecuri ty23/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://w ww.usenix.org/system/files/sec22-xue-diwe n.pdf.
- [29] Yawning Angel. "Obfs4 (the obfourscator)." (Jan. 2019), [Online]. Available: https://github.co m/Yawning/obfs4/blob/master/doc/obfs4-spe c.txt.

Appendices

Proteus Programs Α

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

16

21 22

23

31

38

40

41

42

43 44 45

46

48 49

50 51

52 53

54

58

59 60

61 62

63 64

69

80

81 82

```
1
2
3
4
            /*
* Event handlers
             SET_HANDLER( EV_INIT, evInit );
            SET_HANDLER( EV_NET, evNetRead);
SET_HANDLER( EV_APP, evAppRead);
SET_HANDLER( EV_APP, evAppRead);
    5
    8
             SET HANDLER( *, exitHandler );
 10
11
12
             /*
* Message formats
*/
 13
  14
15
             // encrypted length 
DEFINE EncLenFmt
              { NAME: EncPavloadLen; TYPE: [u8; 2] }
  17
              { NAME: EncPayloadLenTag; TYPE: [u8; 16] };
 18
19
               // the payload
             // Che parson and parson and
 20
24
                * Global variables
*/
            28
29
30
 32
33
34
35
36
37
               * Event handlers */
 39
             fn evInit() {
                        // grab key from environment variable
match getenv<[u8; 32]>("chacha20_key", &global.Key) {
                                    (false, _) => panic(),
                        };
             }
            fn evNetRead() {
    // compute size of the EncLenFmt frame
    let encPayloadLenSize :u64
47
                                     = get_field_size(&EncLenFmt, "EncPayloadLen");
                        let encPayloadLenTagSize :u64
                                      = get_field_size(&EncLenFmt, "EncPayloadLenTag");
 55
56
57
                        let expectedLen :u64 = encPayloadLenSize + encPayloadLenTagSize;
                        let blocking: bool = true;
                         // block until there are at least 'expectedLen' bytes to read
let encLenBytes: [u8; *] = buffer_pop(&RB_net, expectedLen, blocking);
                         let encLen: Fields = match parse(&EncLenFmt, &encLenBytes) {
                                    (true, v) => v,
(false,_) => {
                                              buffer_close_all(); // close connection on error
                                               return;
                        };
 70
71
72
73
74
75
76
77
78
79
                        let encPayloadLen: [u8; 2] = [0u8; 2];
match getField<[u8; *]>(&encLen, "EncPayloadLen", &encPayloadLen) {
                                    false => panic(),
                        3:
                        let encPayloadLenTag: [u8; 16] = [0u8; 16];
match getField<[u8; *]>(&envLen, "EncPayloadLenTag", &encPayloadLenTag) {
                                    false => panic(),
                         };
                       encPavloadLenSize. // message size
&global.InNonce, // nonce and (next line) tag
```

```
87
88
89
                           (true,v) => v,
(false,_) => {
 90
                                 buffer_close_all(); // close connection on error
 91
                                 return:
 92
 93
94
95
96
97
98
99
                    };
              global.InNonce = global.InNonce + 1u64;
               // then, grab the encrypted payload and parse it
let encPayload: [u8; *] = buffer_pop(&RB_net, 1, blocking);
              let payload: Fields = match parse(&EncryptedPayloadFmt, &encPayload) {
  (true,v) => v,
  (false,_) => panic(),
100
101
102
103
104
105
106
107
108
109
              let payload_buffer: BytesMut = get_buffer(65536u64);
let encPayload: [u8; *] = match getField<[u8; *]>(
    &payload_buffer, "EncPayload", spayload_buffer) {
                           false => panic(),
                     };
              let tag_buffer: BytesMut = get_buffer(l6u64);
let encPayloadTag: [u8; *] = match getField([u8; *]>(
    spayload, "EncPayloadTag", stag_buffer) {
      false => panic(),
110
111
112
113
114
115
116
117
118
119
120
121
122
                     };
                  and decrypt it
              &encPayload,
                     &encPayloadSize, &global.InNonce,
123
124
                     &encPayloadTag)
              ł
125
                           (true, v) => v,
125
126
127
128
                           (false
                                       ) =>
                                 buffer_close_all(); // close connection on error
                                return;
129
                          }
130
131
              };
              global.InNonce = global.InNonce + 1u64;
132
133
134
135
              // send the results to the app
buffer_push(&WB_app, &plaintext);
136
137
137
138
139
       fn evAppRead() {
              evappread() {
    // grab data from buffer (from the application)
    let 1 :ul6 = buffer_length(&RB_app);
    let data :[u8; *] = buffer_pop(&RB_app, 1, false);
140
140
141
142
143
144
               // encrypt the length
145
146
147
              let (encLen, encLenTag): ([u8; 2], [u8; 16]) = match encrypt(
    "chacha20-poly1305",
                     &global.Key,
148
                     &l,
2u64,
148
149
150
151
                     &global.OutNonce ) {
                           (true,ciphertext,tag) => (ciphertext, tag),
152
153
154
155
156
                           (false,_,_) => panic(),
              global.OutNonce = global.OutNonce + 1u64;
              // encrypt the payload
let (encPayload, encPayloadTag): ([u8; *], [u8; 16]) = match encrypt(
    "chacha20-poly1305",
150
157
158
159
                    &global.Key,
160
                     &data,
161
162
                     global.OutNonce ) {
163
                           (true,ciphertext,tag) => (ciphertext, tag),
(false,__) => panic(),
164
165
                    };
166
167
              global.OutNonce = global.OutNonce + 1u64;
168
169
               // produce the fram
               let encLenSpec: [u8; *]= match format(&EncLenFmt,
170
170
171
172
173
174
                    $format![
   ("EncPayloadLenTag", encLenTag),
   ("EncPayloadLenTag", encLenTag)
175
176
177
                           (true,v) => v,
                           (false,_) => panic(),
178
179
180
181
              };
               let encPayloadSpec: [u8; *] = match format(&EncryptedPayloadFmt,
                     &format![
                           ("EncPayload", encPayload),
182
183
184
                           ("EncPayloadTag", encPayloadTag),
                    ]
) {
185
185
186
187
188
189
                           (true,v) => v,
(false,_) => panic(),
              };
190
               // send them on the wire
191
192
              buffer_push(&WB_net, &encLenSpec);
buffer_push(&WB_net, &encPayloadSpec);
193
194
```

&encPayloadLenTag) {

```
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
4
        // encrypted length
      DEFINE EnclenFmt
+{ NAME: FixedPreamble; TYPE: [u8; 6] },
 6
7
        { NAME: EncPayloadLen; TYPE: [u8; 2] }
        { NAME: EncPayloadLenTag; TYPE: [u8; 16] };
      00 -45,6 +46,8 00
10
11
12
13
       fn evNetRead() {
            let fixed_preamble_size :u64
             = get_field_size(kEncLenFmt, "FixedPreamble");
// compute size of the EncLenFmt frame
let encPayloadLenSize :u64
14
      +
15
16
17
     = get_field_size(&EncLenFmt, "EncPayloadLen");
@@ -52,7 +55,8 @@
18
              let encPayloadLenTagSize :u64
= get_field_size(&EncLenFmt, "EncPayloadLenTag");
19
20
21
            let expectedLen :u64 = encPayloadLenSize + encPayloadLenTagSize;
let expectedLen :u64 = encPayloadLenSize +
encPayloadLenTagSize + fixed_preamble_size;
22
23
24
25
26
27
             let blocking: bool = true;
      00 -169,6 +173,7 00
28
29
              // produce the frames
30
              let encLenSpec: [u8; *]= match format(&EncLenFmt,
31
32
                  sformat![
  ("FixedPreamble", "123456"),
  ("EncPayloadLen", encLen),
  ("EncPayloadLenTag", encLenTag)
33
34
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:
```

```
 \begin{array}{l} \leftarrow s & (\text{sent out-of-band}) \\ \dots \\ \rightarrow e, es \\ \leftarrow e, ee \end{array}
```

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.

```
* Event handlers
   2
3
4
5
         SET_HANDLER( EV_INIT, evInit );
        SET_HANDLER( EV_NET, evNetRead );
SET_HANDLER( EV_APP, evAppRead );
SET_HANDLER( EV_TIMER, nullHandler );
   6
7
        SET_HANDLER( *, exitHandler );
                                                                  // everything else goes to exitHandler
  10
 11
12
          * Message formats
 13
14
15
16
17
          // corresponds to \rightarrow e, es
         DEFINE Handshakel
         { NAME: InitiatorEphemeralKey; TYPE: [u8; 45] };
 18
19
                               is to <- e, ee
 20
         DEFINE Handshake2
 21
22
23
         { NAME: EncResponderEphemeralKey; TYPE: [u8; 56] },
{ NAME: EncResponderEphemeralKeyTag; TYPE: [u8; 16] };
          // an encrypted message (after handshaking)
 24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
         // an encrypted message (arter name
DEFINE EncryptedPayloadSpec
{ NAME: PayloadSize; TYPE: ul6 },
{ NAME: EncPayload; TYPE: [u8; *] },
          { NAME: EncPayloadTag; TYPE: [u8; 16] };
         GLOBALS {
                let mut CompletedHandshake :bool = false;
let mut OutNonce :u64 = 0u64;
let mut InNonce :u64 = 0u64;
let mut IsInitiator :bool = false;
                let mut ck :[u8;32] = [0u8; 32];
let mut k :[u8;32] = [0u8; 32];
 40
41
42
                // some constants we'll need later
                let byte01 :[u8;1] = [ 0x01; 1 ];
let byte02 :[u8;1] = [ 0x02; 1 ];
 43
44
45
46
47
48
49
                // key material
                let mut EphemeralDHKeyPub :[u8;56] = [0u8; 56];
let mut EphemeralDHKeyPui :[u8;56] = [0u8; 56];
let mut ServerStaticDHKeyPub :[u8;56] = [0u8; 56];
 50
51
52
53
54
55
56
57
58
59
60
61
                let mut StaticDHKeyPub :[u8;56] = [0u8; 56];
let mut StaticDHKeyPri :[u8;56] = [0u8; 56];
           fn evInit() {
               // let's figure out if we're the initiator or responder
// by looking for the is_initiator environment variable
               let mut tmp: [u8;1] = [0u8; 1];
 62
63
64
65
                global.IsInitiator = match getenv<[u8; 1]>("is_initiator", &tmp) {
                       (true, ) => true,
                       (false,_) => false
 66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
                };
                // initialize hash according to InitializeSymmetric(...) func from Noise
let h :[u8;32] = hash(sha256, "Noise_NK_25519_ChaChaPoly_SHA256");
                global.ck = h;
                global.k = h;
                (global.EphemeralDHKeyPub, global.EphemeralDHKeyPri) = genDHKeyPair(56u64);
                match getenv<[u8; 56]>("server_pub", &global.ServerStaticDHKeyPub) {
                       (false,_) => panic(),
                };
                if global.IsInitiator == true {
 81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
                      // compute initial key by first doing DH...
let input_key_material :[u8;56]
                              match DH(global.EphemeralDHKeyPri, global.ServerStaticDHKeyPub) {
                                    (true,v) => v,
(false,_) => panic(),
                      };
                      // update the global key according to Noise HKDF() function
let temp_k :[u8;32] = hmac(sha256, global.ck, input_key_material);
global.ck = hmac(sha256, concatenate(global.ck,global.hyte02));
                } else {
                      // we're the server, grab key pair from environment variable
match getenv<[u8;56]>("dhkey_pub", &global.StaticDHKeyPub) {
                             (false,_) => panic(),
                       match getenv<[u8;56]>("dhkey_pri", &global.StaticDHKeyPri) {
                             (false,_) => panic(),
100
                       }:
100
101
102
103
104
105
         fn evAppRead() {
```

```
if !global.CompletedHandshake {
          before we do anything else, we need to complete the handshake
      if global.IsInitiator {
           // send Handshakel message to responder
let handshakelFields :Fields = create_fields();
           set_field( &handshakelFields,
                "InitiatorEphemeralKey", global.EphemeralDHKeyPub);
            let handshakel :[u8; *] = match format(&Handshakel, &handshakelFields) {
                 (true, v) => v,
                (false,_) => panic(),
           };
          buffer push( &WB net, handshakel );
           // wait for response from responder
let handshake2KeySize :u64
                 = get field size(&Handshake2, "EncResponderEphemeralKey");
           = get_ilera_size,...
let frameContents :[u8; *]
                = match buffer_pop(&RB_net, handshake2KeySize+handshake2TagSize, true) {
                      (true,v) => v,
(false,_) => panic(),
                };
           // parse response (a handshake2 message)
let handshake2 :Fields = match parse(&Handshake2, frameContents) {
                (true,v) => v,
(false,_) => {
    buffer_close_all(); // close connection on error
                      return;
           };
           let encResponderEphemeralKey :BytesMut = get_buffer(64u64);
           if !get_field<[u8; 56]>($handshake2,
    "EncResponderEphemeralKey", &encResponderEphemeralKey) {
    buffer_close_all();
                   return:
           };
           let encResponderEphemeralKeyTag :BytesMut = get_buffer(32u64);
           if !getField<[u8; 16]>(&handshake2,
    "EncResponderEphemeralKeyTag", &encResponderEphemeralKeyTag) {
                  buffer_close_all();
                    return;
           };
           // decrypt to get the responder's ephemeral key
           // delype to get the responder's ephemeral Key
let responderEphemeralKey :[u8; 56] = match decrypt(
    "chacha20-poly1305",
    global.k,
                 &encResponderEphemeralKey,
                handshake2KeySize,
&global.InNonce,
                &encResponderEphemeralKeyTag) {
                      (true,v) => v,
(false,_) => {
    buffer_close_all();
                           return;
                      3
           };
           global.InNonce = global.InNonce + 1u64;
           // update the global key according to Noise HKDF() function
let input_key_material :[u8;56] = responderEphemeralKey;
           let imput_key_material :[05,30] = responsering memoralkey,
let temp_k :[08,32] = hmac(sha256, global.ck, nput_key_material);
global.ck = hmac(sha256, temp_k, global.byte01);
global.k = hmac(sha256, global.ck, global.byte02);
           global.CompletedHandshake = true;
     } else {
           // we're the responder, so we'll wait for the initiator to send Handshakel
let handshakelKeySize :u64
        = get_field_size(sHandshakel, "InitiatorEphemeralKey");
           let frameContents · Bute
                = match buffer_pop(&RB_net, handshakelKeySize, true) {
                     (true,v) => v,
(false,_) => panic(),
                };
            // parse response (a handshakel message)
           let handshakel :Fields = match parse(&Handshakel, frameContents ) {
                (true,v) => v,
(false,_) => {
    buffer_close_all();
                      return;
           };
           let initiatorPK :BytesMut = get_buffer(64u64);
           if !get_field<[u8; 56]>(&Handshakel, "InitiatorEphemeralKey", &initiatorPK) {
                buffer_close_all();
                return;
           // compute our first DH
let input_key_material :[u8;56]
```

```
= match DH(global.StaticDHKeyPri, initiatorPK) {
                           (true,v) => v,
(false,_) => panic(),
                let temp k :[u8:32] = hmac(sha256, global.ck, input key material);
                global.ck = hmac(sha256, temp_k, global.byte01);
global.k = hmac(sha256, global.ck, global.byte02);
                &global.k.
                      &global.EphemeralDHKeyPub,
                      56u64,
                     &global.OutNonce ) {
                          (true,ciphertext,tag) => (ciphertext, tag),
(false,__) => panic(),
                };
                global.OutNonce = global.OutNonce + 1u64;
                let handshake2Fields :Fields = create_fields();
                set_field( &handshake2Pields,
   "EncResponderEphemeralKey", encPayload );
set_field( &handshake2Fields,
                      "EncResponderEphemeralKeyTag", encPayloadTag );
                let handshake2 :Bytes = match format(
    &Handshake2, &handshake2Fields ) {
      (true, v) => v,
                           (false,_) => panic(),
                };
                buffer push( &WB net, handshake2 );
                // and compute our final global key
let input_key_material2 :[u8;56]
                     = match DH(global.EphemeralDHKevPri, initiatorPK) {
                           (true,v) => v,
(false,_) => panic(),
                /,
let temp_k :[u8;32] = hmac(sha256,global.ck,input_key_material2);
                global.k = hmac(sha256, temp_k,global.byte01);
global.k = hmac(sha256, global.ck, global.byte02);
                global.CompletedHandshake = true;
     } else {
           // handshake completed, so send data in AEAD
           // get data from app
let 1 :u64 = buffer_length(&RB_app);
           let dat:Bytes = match buffer_pop(&RB_app, 1, true) {
  (true,v) => v,
  (false,_) => panic(),
           };
           // encrypt the payload
let (encPayload, encPayloadTag): ([u8; *], [u8; 16]) = match encrypt(
                 chacha20-poly1305",
                &αlobal.k.
                &data,
                &global.OutNonce ) {
                     (true,ciphertext,tag) => (ciphertext, tag),
(false,__) => panic(),
           global.OutNonce = global.OutNonce + 1u64;
           // put the ciphertext in its frame
let encryptedPayloadFields :Fields = create_fields();
           let entrypeterayloadrens :rens - clear_inens(),
set_field(sencryptedPayloadrends, "Payloadsire", 1);
set_field(sencryptedPayloadFields, "EncPayload", encPayload);
set_field(sencryptedPayloadFields, "EncPayloadTag", encPayloadTag);
           let encPayloadSpec :Bytes
                = match format(&EncryptedPayloadSpec, &encryptedPayloadFields) {
                     (true,v) => v,
(false,_) => panic(),
           };
          // and send it!
buffer_push( &WB_net, encPayloadSpec );
fn evNetRead() {
     if !global.CompletedHandshake {
             / before we do anything else, we need to complete the handshake
           if global.IsInitiator {
                // send Handshakel message to responder
let handshakelFields :Fields = create_fields();
                let manusankerleds :reids : reids_reids(),
set_field (shandshakelFields,
    "InitiatorEphemeralKey", global.EphemeralDHKeyPub);
let handshakel :Bytes = match format(&Handshakel, &handshakelFields) {
                      (true, v) \Rightarrow v,
                      (false,_) => panic(),
                buffer_push( &WB_net, handshakel );
                // wait for response from responder
                let handshake2TagSize :u64
                      = get_field_size(&Handshake2,"EncResponderEphemeralKeyTag");
```

```
322
 323
324
325
 326
327
 328
329
 330
 331
332
333
 334
 335
336
337
 338
339
340
341
 342
343
344
345
 346
347
348
 349
 350
351
352
 353
354
355
 356
357
 358
359
 360
 361
362
363
364
 365
366
367
 368
369
370
371
 372
373
374
375
 376
377
378
379
 380
381
382
 383
 384
385
 386
 387
 388
389
 390
391
 392
393
 394
 395
 396
397
 398
 399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
 420
420
421
422
423
424
425
426
427
428
```

};

global.OutNonce = global.OutNonce + 1u64;

```
let frameContents :Bytes
                                                                                                                                                 429
             riamecontents to ress
= match buffer pop(&B_net, handshake2KeySize+handshake2TagSize, true) {
  (true,v) => v,
  (false,_) => panic(),
                                                                                                                                                430
431
432
433
            1:
                                                                                                                                                434
435
436
                                   (a handshake2 messag
      let handshake2 :Fields = match parse(&Handshake2, frameContents) {
             (true,v) => v,
(false,_) => {
                                                                                                                                                 437
                                                                                                                                                437
438
439
440
441
                   buffer_close_all(); // close connection on error
                   return;
      };
                                                                                                                                                 442
443
                                                                                                                                                 444
      let encResponderEphemeralKey :BytesMut = get_buffer(64u64);
                                                                                                                                                445
446
447
448
449
      if !get_field<[u8; 56]>(&handshake2,
    "EncResponderEphemeralKey", &encResponderEphemeralKey) {
            buffer_close_all();
              return;
                                                                                                                                                 450
451
      };
      let encResponderEphemeralKeyTag :BytesMut = get buffer(32u64);
                                                                                                                                                 452
                                                                                                                                                453
454
455
456
457
458
459
       if !getField<[u8; 16]>(&handshake2,
              "EncResponderEphemeralKeyTag", &encResponderEphemeralKeyTag) {
            buffer_close_all();
             return;
      };
       // decrypt to get the responder's ephemeral key
let responderEphemeralKey: [u8; *] = match decrypt(
    "chacha20-poly1305",
    "let")
                                                                                                                                                460
461
462
463
464
            global.k,
&encResponderEphemeralKey,
handshake2KeySize,
&global.InNonce,
                                                                                                                                                465
466
             &encResponderEphemeralKeyTag) {
                                                                                                                                                 467
                   (true,v) => v,
(false,_) => {
    buffer_close_all();
                                                                                                                                                468
469
470
471
                         return;
                                                                                                                                                471
472
473
474
475
      };
       global.InNonce = global.InNonce + 1u64;
                                                                                                                                                475
476
477
478
479
      // update the global key according to Noise HKDF() function
let input_key_material :[u8;56] = responderEphemeralKey;
let temp_k :[u8;32] = hmac(sha256, global.ck, input_key_material);
global.ck = hmac(sha256, temp_k, global.byte01);
global.k = hmac(sha256, global.ck, global.byte02);
                                                                                                                                                 480
481
                                                                                                                                                 482
                                                                                                                                                483
484
485
      global.CompletedHandshake = true;
} else {
                                                                                                                                                 486
      // we're the responder, so we'll wait for the initiator to send Handshakel {\tt let\ handshakelKeySize\ :}u64
                                                                                                                                                487
488
489
             = get_field_size(&Handshakel, "InitiatorEphemeralKey");
                                                                                                                                                 490
                                                                                                                                                490
491
492
493
       let frameContents :Bytes
    = match buffer_pop(&RB_net, handshakelKeySize, true) {
                   (true, v) => v
                    (false,_) => panic(),
                                                                                                                                                 494
                                                                                                                                                495
496
            };
      // parse response (a handshakel message) let handshakel :Fields = match parse(&Handshakel, frameContents ) { (true, v) \Rightarrow v, (false,_) => {
                                                                                                                                                 497
                                                                                                                                                 498
                                                                                                                                                499
500
                  buffer close all();
                                                                                                                                                 501
                                                                                                                                                 502
                   return;
                                                                                                                                                502
503
504
      };
                                                                                                                                                 505
                                                                                                                                                506
507
      let initiatorPK :BytesMut = get_buffer(64u64);
      if !get_field<[u8; 56]>(&Handshakel,"InitiatorEphemeralKey", &initiatorPK) {
                                                                                                                                                 508
             buffer_close_all();
                                                                                                                                                 509
                                                                                                                                                510
511
512
             return;
      };
                                                                                                                                                512
513
514
515
516
       let input_key_material :[u8;56]
    = match DH(global.StaticDHKeyPri, initiatorPK) {
                   (true, v) => v,
                    (false,_) => panic(),
                                                                                                                                                 517
                                                                                                                                                518
519
       };
      let temp_k :[u8; 32] = hmac(sha256, global.ck, input_key_material);
global.ck = hmac(sha256, temp_k, global.byte01 );
global.k = hmac(sha256, global.ck, global.byte02 );
       // send our ephemeral key, encrypted, to the initiator
let (encPayload, encPayloadTag): ([u8; *], [u8; 16]) = match encrypt(
    "chacha20-poly1305",
             &global.k,
            &global.EphemeralDHKeyPub,
56u64,
&global.OutNonce ) {
                    (true,ciphertext,tag) => (ciphertext, tag),
                   (false,_,_) => panic(),
```

```
let handshake2Fields :Fields = create_fields();
set_field( &handshake2Fields,
    "EncResponderEphemeralKey", encPayload );
               set_field( &handshake2Fields,
    "EncResponderEphemeralKeyTag", encPayloadTag );
               let handshake2 :Bytes = match format(
                    &Handshake2, &handshake2Fields )
  (true,v) => v,
  (false,_) => panic(),
               buffer_push( &WB_net, handshake2 );
                // and compute our final global key
               (true, v) => v,
                         (false,_) => panic(),
               //,
let temp_k :[U8;32] = hmac(sha256,global.ck,input_key_material2);
global.ck = hmac(sha256,temp_k,global.byte01);
               global.k = hmac(sha256, global.ck, global.byte02);
               global.CompletedHandshake = true;
     ) else (
          // handshake completed, so grab the data off of the wire
let payloadSizeLen :u64
                = get_field_size(&EncryptedPayloadSpec,"PayloadSize");
          let encPayloadTagLen :u64
    = get_field_size(&EncryptedPayloadSpec,"EncPayloadTag");
          let 1 :ul6 = buffer peek(&RB net, payloadSizeLen);
          let frameContents :Bytes = match buffer_pop(
               RB_net,
               payloadSizeLen + 1 + encPayloadTagLen,
               true) {
    (true,v) => v,
                    (false,_) => panic(),
               };
          let encryptedPayload: Fields = match parse(&EncryptedPayloadSpec, &frameContents ) {
               (true,v) => v,
(false,_) => {
                    buffer_close_all(); // close connection on error
                    return:
          iet encPayload: BytesMut = get_buffer(65535u64);
          if !get_field<[u8; *]>(&EncryptedPayloadSpec, "EncPayload", &encPayload )
              panic();
          };
          let encPayloadTag: BytesMut = get_buffer(32u64);
if !getField<[u8; 32]>(&EncryptedPayloadSpec, "EncPayloadTag", &encPayloadTag ) {
               panic();
          }:
          // decrypt returns (num_bytes,val) tuple
let plaintext :[u8; *] = match decrypt(
    "chacha20-poly1305",
    &global.k,
               &encPayload,
               &global.InNonce.
               &encPayloadTag)
                    (true, v) => v,
                    (false, ) => {
                         buffer_close_all(); // close connection on error
                        return;
                    }
          global.InNonce = global.InNonce + 1u64;
            / send the results to the app
          buffer_push( &WB_app, plaintext );
}
fn nullHandler() {
fn exitHandler() {
     exit(0u32);
}
```

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.

Functionality **B.1**

Functionality	106	
SECTION - I/O RELATED	107 108 109	buffer_close_all - closes all con SYNOPSIS
	110 111	<pre>fn buffer_close_all();</pre>
NAME	112	DESCRIPTION
buffer_length - get number of bytes available in a buffer	114	Equivalent to calling:
SYNOPSIS	116	in in
<pre>fn buffer_length(b: &Buffer) -> usize;</pre>	118	<pre>buffer_close(rb_app); buffer_close(rb_apt);</pre>
DESCRIPTION	120	<pre>buffer_close(rb_net); buffer_close(wb_app); buffer_close(wb_apt);</pre>
Gets the number of bytes present in a buffer.	121 122 123	<pre>buffer_close(wb_net);</pre>
RETURN VALUE		RETURN VALUE
The number of bytes present in the buffer. No error value is specified.	123 126 127	N/A
	128	
NAME	129 130	
buffer_peek - get a copy of n bytes from a buffer without removing them	131 132	SECTION - UTILITY FUNCTIONS
SYNOPSIS		NAME
fn buffer_peek(b: &ReadBuffer, n: usize) -> (bool, Bytes);	135 136	getenv – gets an environment vari
DESCRIPTION		SYNOPSIS
Gets the first n bytes of data present in the buffer. The buffer is not	139 140	fn getenv <t>(name: &str, value: &</t>
modified as a result of this operation.		DESCRIPTION
RETURN VALUE	143 144	Gets en environment variable of t
buffer_peek() returns a pair of values. The first element of the pair indicates if the full peek could be performed (true if so, otherwise the value	145 146	argument if the variable is defin
is false). The second element contains the copied data from the buffer (with length equal to the minimum of n and buffer_length(b)).	147 148	RETURN VALUE
	149 150	Returns true if the environment v argument; false otherwise.
NAME	151 152	
buffer_pop - removes bytes from a buffer	153 154	NAME
SYNOPSIS	155 156	get_random_bytes - generates a nu
<pre>fn buffer_pop(b: &mut ReadBuffer, n: usize, blocking: bool) -> (bool, Bytes);</pre>	157 158	(not suitable
DESCRIPTION	159 160	SYNOPSIS
Removes the first n bytes from buffer b. If there are fewer than n bytes	161 162	<pre>fn get_random_bytes(n: usize) -> </pre>
available, this function does not modify the buffer. If the blocking parameter is set to true, this function blocks until there is n bytes of data in the	163 164	DESCRIPTION
buffer.	165 166	Generates n uniformly random byte necessarily cryptographically str
RETURN VALUE		RETURN VALUE
For the first return value, returns true if the pop was successful (i.e., at least n bytes were removed from the buffer); false otherwise. The second	169 170	The n randomly sampled bytes.
argument returns the bytes that were removed.	171 172	
		NAME
NAME	175 176	log – logs a string
buffer_push - adds bytes to a buffer.		SYNOPSIS
SYNOPSIS	179 180	fn log(line: &str);
<pre>fn buffer_push(b: &mut WriteBuffer, data: Bytes) -> bool;</pre>	181 182	DESCRIPTION
DESCRIPTION	183 184	Writes the input line out to the
Adds the input data to the buffer, if there is enough room in the buffer. Otherwise, the buffer is not modified.	185 186	RETURN VALUE
RETURN VALUE	187 188	N/A
True if the data was successfully added to the buffer; false otherwise.	189 190	
	191 192	NAME
NAME	193 194	panic - exits the program due to
buffer_close - close the connection associated with a buffer.	195 196	SYNOPSIS
SYNOPSIS	197 198	fn panic();
<pre>fn buffer_close(b: &mut Buffer);</pre>	199 200	DESCRIPTION
DESCRIPTION	201 202	Closes the program and network co
Closes the connection associated with the given buffer. This operation is	203 204	RETURN VALUE
analogous to calling close on a buffer.	205 206	N/A
RETURN VALUE	207 208	
N/A	209 210	NAME
	211 212	exit - exits the program cleanly
NAME	213	· · · · · · · · · · · · · · · · · · ·

<pre>buffer_close_all - closes all connections</pre>	
SYNOPSIS	
<pre>fn buffer_close_all();</pre>	
DESCRIPTION	
Equivalent to calling:	
<pre>buffer_close(rb_app); buffer_close(rb_net); buffer_close(wb_app); buffer_close(wb_net);</pre>	
RETURN VALUE	
N/A	
SECTION - UTILITY FUNCTIONS	
NAME	
getenv - gets an environment variable	
SYNOPSIS	
<pre>fn getenv<t>(name: &str, value: &mut T) -> bool;</t></pre>	
DESCRIPTION	
Gets en environment variable of type T. The value is stored in the 'val argument if the variable is defined and can be cast to type T.	ue'
RETURN VALUE	
Returns true if the environment variable was successfully stored in the argument; false otherwise.	value
NAME	
get_random_bytes - generates a number of random bytes (not suitable for cryptographic use)	
SYNOPSIS	
<pre>fn get_random_bytes(n: usize) -> Bytes;</pre>	
DESCRIPTION	
Generates n uniformly random bytes and returns them. The bytes are not necessarily cryptographically strong.	
RETURN VALUE	
The n randomly sampled bytes.	
NAME	
log - logs a string	
SYNOPSIS	
<pre>fn loq(line: &str);</pre>	
DESCRIPTION	
Writes the input line out to the system log (defined as stderr).	
RETURN VALUE	
N/A	
NAME	
panic - exits the program due to an error	
SYNOPSIS	
<pre>fn panic();</pre>	
DESCRIPTION	
Closes the program and network connections associated with the program $% \left({{{\boldsymbol{x}}_{i}}} \right)$	
Closes the program and network connections associated with the program \ensuremath{RETURN} VALUE	

214 215	SYNOPSIS	322 323	
216 217	<pre>fn exit(exit_code: u32);</pre>	324 325	SECTION - PROTOCOL MESSAGE MANIPULATION
218	DESCRIPTION	326	
219 220 221	Closes the program and network connections associated with the program, returning the specified exit code.	327 328 329	NAME create_fields - create a new empty field object
222 223	RETURN VALUE	330 331	SYNOPSIS
224		332	
225 226	N/A	333 334	<pre>create_fields() -> Fields;</pre>
227 228		335 336	DESCRIPTION
229 230	NAME	337 338	Creates an initialized, empty Fields object. The Fields object is used to store and retrieve message field values by name for message formatting.
231	arm_timer - adds a timer event to the queue	339	
232 233	SYNOPSIS	340 341	RETURN VALUE
234 235	<pre>fn arm_timer(k: u8, t: usize);</pre>	342 343	A new Fields object.
236 237	DESCRIPTION	344 345	
238 239		346 347	NAME
240	Creates a timer with number k that expires t seconds from the calling time. If called with a number of a timer that was already set, this function resets	348	set_field - set a field value
241 242	the timer.	349 350	SYNOPSIS
243 244	RETURN VALUE	351 352	<pre>set_field<t>(fields: &mut Fields, name : &str, value: T);</t></pre>
245 246	N/A	353 354	DESCRIPTION
247 248		355	
249	NAME	357	Sets the field with the given name to by of type T and the given value.
250 251	disarm_timer - disables a set timer	358 359	RETURN VALUE
252 253	SYNOPSIS	360 361	N/A
254 255	<pre>fn disarm_timer(k: u8);</pre>	362 363	
256		364	NAME
257 258	DESCRIPTION	365 366	get_field - get a field value
259 260	Disables the kth timer, if it was previously armed. Otherwise, this function has no effect.	367 368	SYNOPSIS
261 262	RETURN VALUE	369 370	<pre>get_field<t>(fields: &Fields, name: &str, value: &mut T) -> bool;</t></pre>
263 264	N/A	371 372	DESCRIPTION
265 266		373 374	Gets the value of the field with the specified name and type and stores it in
267 268	NAME	375 376	the value argument. Fails on type mismatch or if the name was not set.
269		377	RETURN VALUE
270 271	get_timer - get the number of the last timer that expired	378 379	true if the value fetch was successful; false otherwise.
272 273	SYNOPSIS	380 381	
274 275	<pre>fn get_timer() -> u8;</pre>	382 383	NAME
276 277	DESCRIPTION	384 385	get_field_size - gets the size of a field in a message format.
278 279	On a timer expired event, this function can be used to get the number of the timer that expired.	386 387	SYNOPSIS
280 281	RETURN VALUE	388 389	<pre>get_field_size<t>(format: &MsgFormat, name: &str) -> usize;</t></pre>
282		390	
283 284	The number of the last timer that expired.	392	DESCRIPTION
285 286		393 394	Returns the statically-defined size of a field with the given name.
287 288	NAME	395 396	RETURN VALUE
289 290	concatenate - join two byte objects together	397 398	The size of the field, or 0 if the field was not present or defined to have variable size.
291 292	SYNOPSIS	399 400	
293	<pre>fn concatenate(bl: &Bytes, b2: &Bytes) -> Bytes;</pre>	401 402	
294 295	DESCRIPTION	403	NAME
296 297	The output is a copy of bl and b2 joined together in sequence.	404 405	format - try to create a formatted byte string
298 299	RETURN VALUE	406 407	SYNOPSIS
300 301	b1 b2	408 409	<pre>format(format: &MsgFormat, fields: &Fields) -> (bool, Bytes);</pre>
302 303		410 411	DESCRIPTION
304 305	NAME	412 413	Attempts to format a byte string according to the specified message format and included field values.
306 307	get_buffer - gets a new mutable buffer	414 415	RETURN VALUE
308 309	get_builer gets a new mutable builer	416 417	(true, b) for a formatted byte string b if formatting was successful.
310 311		418 419	(false, empty bytes) if formatting was not successful.
312	<pre>fn get_buffer(capacity: usize) -> BytesMut;</pre>	420	
313 314	DESCRIPTION	421 422	NAME
315 316	Returns an empty BytesMut container with the specifier capacity.	423 424	parse - try to parse a byte string into the specified fields.
317 318	RETURN VALUE	425 426	SYNOPSIS
319 320 321	An empty BytesMut container.	427 428	parse(format: &MsgFormat, data: &bytes) -> (bool, Fields);
321		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;
          false otherwise. If the parse was successful, the returned Fields object will contain the parsed fields. Otherwise, it will be empty.
438
439
441
442
443
       SECTION - CRYPTOGRAPHIC FUNCTIONS
445
       We assume that ProtoSpec supports a standard set of cryptographic functionalities, for example, those specified in the RustCrypto library <https://github.com/RustCrypto>.
446
447
```

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.

С **Proteus Grammar**

29

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) * ~ "*/" }

            identifier = 0{ (("_"\ASCII_ALPHA)~("_"\ASCII_ALPHANUMERIC)") }
dotted_identifier = { identifier ~ ("." ~ identifier) }
compound_identifier = { "(" ~ dotted_identifier ~ ("," ~ dotted_identifier) + ~ ")" }
basic_or_compound_identifier = { dotted_identifier | compound_identifier }
    5
     6
             numeric_type = { "u8" | "u16" | "u32" | "u64" | "i8" | "i16" | "i32" | "i64" }
  10
           numeric_type = { "us" | "ul6" | "u32" | "u64" | "i6" | "i6" | "i32"
basic_type = { numeric_type | "bool" | "char" }
concrete_array_type = { "[" ~ type ~ ";" ~ numeric_literal ~ "]" }
dynamic_array_type = { "[" ~ type ~ ";" ~ "** ~ "]" }
flexible_array_type = { "[" ~ type ~ ";" ~ disj ~ "]" }
array_type = { concrete_array_type | dynamic_array_type }
custom_type = { identifier }
une = { being type is your type | units }
  13
  16
            custom_type = { remtring }
type = { basic_type | array_type | custom_type }
compound_type = { "(" ~ type ~ ("," ~ type)+ ~ ")'
basic_or_compound_type = { type | compound_type }
  17
  20
  21
22
23
           template_type = { "<" ~ basic_or_compound_type ~ ("," ~ basic_or_compound_type)* ~ ">" }
            numeric_literal = @{ ASCII_DIGIT+ }
hex_literal = @{ "Ox" ~ ASCII_HEX_DIGIT+ }
  24
  25
             typed_numeric_literal = ${ hex_literal | (numeric_literal~numeric_type) }
           string literal = ${ "\"" ~ inner ~ "\"" }
inner = @{ char* }
  27
28
             inner = 0
char = {
```

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

Listing 8: Parsing expression grammar recognizing the Proteus language.