

# Combining Model Learning and Model Checking to Analyze TCP Implementations

Paul Fiterău-Broștean, Ramon Janssen, and Frits Vaandrager<sup>(✉)</sup>

Institute for Computing and Information Sciences,  
Radboud University, P.O. Box 9010, 6500 GL Nijmegen, The Netherlands  
F.Vaandrager@cs.ru.nl

**Abstract.** We combine model learning and model checking in a challenging case study involving Linux, Windows and FreeBSD implementations of TCP. We use model learning to infer models of different software components and then apply model checking to fully explore what may happen when these components (e.g. a Linux client and a Windows server) interact. Our analysis reveals several instances in which TCP implementations do not conform to their RFC specifications.

## 1 Introduction

Our society has become completely dependent on network and security protocols such as TCP/IP, SSH, TLS, Bluetooth, and EMV. Protocol specification or implementation errors may lead to security breaches or even complete network failures, and hence many studies have applied model checking to these protocols in order to find such errors. Since exhaustive model checking of protocol implementations is usually not feasible [24], two alternative approaches have been pursued in the literature. This article proposes a third approach.

A first approach, followed in many studies, is to use model checking for analysis of models that have been handcrafted starting from protocol standards. Through this approach many bugs have been detected, see e.g. [8, 11, 21, 26, 28, 45]. However, as observed in [10], the relationships between a handcrafted model of a protocol and the corresponding standard are typically obscure, undermining the reliability and relevance of the obtained verification results. In addition, implementations of protocols frequently do not conform to their specification. Bugs specific to an implementation can never be captured using this way of model checking. In [18], for instance, we showed that both the Windows 8 and Ubuntu 13.10 implementations of TCP violate the standard. In [40], new security flaws were found in three of the TLS implementations that were analyzed, all due

---

P. Fiterău-Broștean—Supported by NWO project 612.001.216, Active Learning of Security Protocols (ALSEP).

R. Janssen—Supported by STW project 13859, Supersizing Model-Based Testing (SUMBAT).

F. Vaandrager—Supported by STW project 11763, Integrating Testing And Learning of Interface Automata (ITALIA).

to violations of the standard. In [13, 46] it was shown that implementations of a protocol for Internet banking and of SSH, respectively, violate their specification.

A second approach has been pioneered by Musuvathi and Engler [31]. Using the CMC model checker [32], they model checked the “hardest thing [they] could think of”, the Linux kernel’s implementation of TCP. Their idea was to run the *entire* Linux kernel as a CMC process. Transitions in the model checker correspond to events like calls from the upper layer, and sending and receiving packets. Each state of the resulting CMC model is around 250 kilobytes. Since CMC cannot exhaustively explore the state space, it focuses on exploring states that are the most different from previously explored states using various heuristics and by exploiting symmetries. Through their analysis, Musuvathi and Engler found four bugs in the Linux TCP implementation. One could argue that, according to textbook definitions of model checking [7, 17], what Musuvathi and Engler do is not model checking but rather a smart form of testing.

The approach we explore in this paper uses model learning. Model learning, or active automata learning [1, 6, 44], is emerging as a highly effective technique to obtain models of protocol implementations. In fact, all the standard violations reported in [13, 18, 40, 46] have been discovered (or reconfirmed) with the help of model learning. The goal of model learning is to obtain a state model of a black-box system by providing inputs to and observing outputs. This approach makes it possible to obtain models that fully correspond to the observed behavior of the implementation. Since the models are derived from a finite number of observations, we can (without additional assumptions) never be sure that they are correct: even when a model is consistent with all the observations up until today, we cannot exclude the possibility that there will be a discrepancy tomorrow. Nevertheless, through application of conformance testing algorithms [27], we may increase confidence in the correctness of the learned models. In many recent studies, state-of-the-art tools such as LearnLib [44] routinely succeeded to learn correct models efficiently. In the absence of a tractable white-box model of a protocol implementation, a learned model is often an excellent alternative that may be obtained at relatively low cost.

The main contribution of this paper is the combined application of model checking, model learning and abstraction techniques in a challenging case study involving Linux, Windows and FreeBSD implementations of TCP. Using model learning and abstraction we infer models of different software components and then apply model checking to explore what may happen when these components (e.g. a Linux client and a Windows server) interact.

The idea to combine model checking and model learning was pioneered in [36], under the name of *black box checking*. In [30], a similar methodology was introduced to use learning and model checking to obtain a strong model-based testing approach. Following [30, 36], model checkers are commonly used to analyze models obtained via automata learning. However, most of these applications only consider specifications of a single system component, and do not analyze networks of learned models. An exception is the work of Shahbaz and Groz [42] on integration testing, in which learned component models are composed and then analyzed using reachability analysis in order to find integration faults. Our results

considerably extend our previous work on learning fragments of TCP [18] since we have (1) added inputs corresponding to calls from the upper layer, (2) added transmission of data, (3) inferred models of TCP clients in addition to servers, and (4) learned models for FreeBSD in addition to Windows and Linux. Abstraction is the key for scaling existing automata learning methods to realistic applications. In order to obtain tractable models we use the theory of abstractions from [2], which in turn is inspired by earlier work on predicate abstraction [16, 29]. Our use of abstractions is similar to that of Cho et al. [14], who used abstractions to infer models of realistic botnet command and control protocols. Whereas in our previous studies on model learning the abstractions were implemented by ad-hoc Java programs, we now define them in a more systematic manner. We provide a language for defining abstractions, and from this definition we automatically generate mapper components for learning and model checking.

Our method may be viewed as a smart black-box testing approach that combines the strengths of model learning and model checking. The main advantage of our method compared to approaches in which models are handcrafted based on specifications is that we analyze the “real thing” and may find “bugs” in implementations. In fact, our analysis revealed several instances in which TCP implementations do not conform to the standard. Compared to the white-box approach of Musuvathi and Engler [31], our black-box method has several advantages. First of all, we obtain explicit component models that can be fully explored using model checking. Also, our method appears to be easier to apply and is more flexible. For instance, once we had learned a model of the Linux implementation it took just two days to learn a model of the Windows implementation. In the approach of [31], one first would need to get access to the proprietary code from Microsoft, and then start more or less from scratch from an entirely different code base. In contrast, using our approach it is possible to learn a model of any TCP implementation within a few days. Besides these practical benefits, there is also an important philosophical advantage. If one constructs a model of some real-world phenomenon or system and makes claims based on this model then, in line with Popper [37], we think this model ought to be falsifiable. Our model of the Windows8 TCP client is included in the paper in Fig. 2, and all Mealy machine and nuSMV models are available at <http://www.sws.cs.ru.nl/publications/papers/fvaan/FJV16/>. Our notion of state is clear and based on the Nerode congruence [33]: two traces lead to the same state unless there is a distinguishing suffix. Any researcher can study our models and point out mistakes. In contrast, the model of Musuvathi is specified implicitly through heuristics (when have we seen a state before?) that are programmed on top of the Linux implementation. As a result, falsification of their model is virtually impossible.

## 2 Background on Model Learning

**Mealy Machines.** During model learning, we represent protocol entities as Mealy machines. A *Mealy machine* is a tuple  $\mathcal{M} = \langle I, O, Q, q^0, \rightarrow \rangle$ , where  $I, O,$

and  $Q$  are finite sets of *input actions*, *output actions*, and *states*, respectively,  $q^0 \in Q$  is the *initial state*, and  $\rightarrow \subseteq Q \times I \times O \times Q$  is the *transition relation*. We write  $q \xrightarrow{i/o} q'$  if  $(q, i, o, q') \in \rightarrow$ . We assume  $\mathcal{M}$  to be *input enabled* (or *completely specified*) in the sense that, for each state  $q$  and input  $i$ , there is a transition  $q \xrightarrow{i/o} q'$ , for some  $o$  and  $q'$ . We call  $\mathcal{M}$  *deterministic* if for each state  $q$  and input  $i$  there is exactly one output  $o$  and one state  $q'$  such that  $q \xrightarrow{i/o} q'$ . We call  $\mathcal{M}$  *weakly deterministic* if for each state  $q$ , input  $i$  and output  $o$  there is exactly one state  $q'$  with  $q \xrightarrow{i/o} q'$ .

Let  $\sigma = i_1 \cdots i_n \in I^*$  and  $\rho = o_1 \cdots o_n \in O^*$ . Then  $\rho$  is an *observation* triggered by  $\sigma$  in  $\mathcal{M}$ , notation  $\rho \in A_{\mathcal{M}}(\sigma)$ , if there are  $q_0 \cdots q_n \in Q^*$  such that  $q_0 = q^0$  and  $q_{j-1} \xrightarrow{i_j/o_j} q_j$ , for all  $j$  with  $0 \leq j < n$ . If  $\mathcal{M}$  and  $\mathcal{M}'$  are Mealy machines with the same inputs  $I$  and outputs  $O$ , then we write  $\mathcal{M} \leq \mathcal{M}'$  if, for each  $\sigma \in I^*$ ,  $A_{\mathcal{M}}(\sigma) \subseteq A_{\mathcal{M}'}(\sigma)$ . We say that  $\mathcal{M}$  and  $\mathcal{M}'$  are (*behaviorally*) *equivalent*, notation  $\mathcal{M} \approx \mathcal{M}'$ , if both  $\mathcal{M} \leq \mathcal{M}'$  and  $\mathcal{M}' \leq \mathcal{M}$ .

If  $\mathcal{M}$  is deterministic, then  $A_{\mathcal{M}}(\sigma)$  is a singleton set for each input sequence  $\sigma$ . In this case,  $\mathcal{M}$  can equivalently be represented as a structure  $\langle I, O, Q, q^0, \delta, \lambda \rangle$ , with  $\delta : Q \times I \rightarrow Q$ ,  $\lambda : Q \times I \rightarrow O$ , and  $q \xrightarrow{i/o} q' \Rightarrow \delta(q, i) = q' \wedge \lambda(q, i) = o$ .

**MAT Framework.** The most efficient algorithms for model learning all follow the pattern of a *minimally adequate teacher (MAT)* as proposed by Angluin [6]. In the MAT framework, learning is viewed as a game in which a learner has to infer an unknown automaton by asking queries to a teacher. The teacher knows the automaton, which in our setting is a deterministic Mealy machine  $\mathcal{M}$ . Initially, the learner only knows the inputs  $I$  and outputs  $O$  of  $\mathcal{M}$ . The task of the learner is to learn  $\mathcal{M}$  through two types of queries:

- With a *membership query*, the learner asks what the response is to an input sequence  $\sigma \in I^*$ . The teacher answers with the output sequence in  $A_{\mathcal{M}}(\sigma)$ .
- With an *equivalence query*, the learner asks whether a hypothesized Mealy machine  $\mathcal{H}$  is correct, that is, whether  $\mathcal{H} \approx \mathcal{M}$ . The teacher answers *yes* if this is the case. Otherwise it answers *no* and supplies a *counterexample*, which is a sequence  $\sigma \in I^*$  that triggers a different output sequence for both Mealy machines, that is,  $A_{\mathcal{H}}(\sigma) \neq A_{\mathcal{M}}(\sigma)$ .

Starting from Angluin’s seminal  $L^*$  algorithm [6], many algorithms have been proposed for learning finite, deterministic Mealy machines via a finite number of queries. We refer to [23] for recent overview. In applications in which one wants to learn a model of a black-box reactive system, the teacher typically consists of a System Under Learning (SUL) that answers the membership queries, and a conformance testing tool [27] that approximates the equivalence queries using a set of *test queries*. A test query consists of asking to the SUL for the response to an input sequence  $\sigma \in I^*$ , similar to a membership query.

**Abstraction.** We recall relevant parts of the theory of abstractions from [2]. Existing model learning algorithms are only effective when applied to Mealy machines with small sets of inputs, e.g. fewer than 100 elements. Practical systems like TCP, however, typically have huge alphabets, since inputs and outputs carry parameters of type integer or string. In order to learn an over-approximation of a “large” Mealy machine  $\mathcal{M}$ , we place a transducer in between the teacher and the learner, which translates concrete inputs in  $I$  to abstract inputs in  $X$ , concrete outputs in  $O$  to abstract outputs in  $Y$ , and vice versa. This allows us to abstract a Mealy machine with concrete symbols in  $I$  and  $O$  to a Mealy machine with abstract symbols in  $X$  and  $Y$ , reducing the task of the learner to inferring a “small” abstract Mealy machine.

Formally, a *mapper* for inputs  $I$  and outputs  $O$  is a deterministic Mealy machine  $\mathcal{A} = \langle I \cup O, X \cup Y, R, r_0, \delta, \lambda \rangle$ , where  $I$  and  $O$  are disjoint sets of *concrete input and output symbols*,  $X$  and  $Y$  are disjoint sets of *abstract input and output symbols*, and  $\lambda : R \times (I \cup O) \rightarrow (X \cup Y)$ , the *abstraction function*, respects inputs and outputs, that is, for all  $a \in I \cup O$  and  $r \in R$ ,  $a \in I \Leftrightarrow \lambda(r, a) \in X$ .

Basically, the *abstraction* of Mealy machine  $\mathcal{M}$  via mapper  $\mathcal{A}$  is the Cartesian product of the underlying transition systems. Let  $\mathcal{M} = \langle I, O, Q, q_0, \rightarrow_{\mathcal{M}} \rangle$  be a Mealy machine and let  $\mathcal{A} = \langle I \cup O, X \cup Y, R, r_0, \delta_{\mathcal{A}}, \lambda_{\mathcal{A}} \rangle$  be a mapper. Then  $\alpha_{\mathcal{A}}(\mathcal{M})$ , the *abstraction of  $\mathcal{M}$  via  $\mathcal{A}$* , is the Mealy machine  $\langle X, Y \cup \{\perp\}, Q \times R, (q_0, r_0), \rightarrow_{\alpha} \rangle$ , where  $\perp \notin Y$  is a fresh output and  $\rightarrow_{\alpha}$  is given by the rules

$$\frac{q \xrightarrow{i/o}_{\mathcal{M}} q', r \xrightarrow{i/x}_{\mathcal{A}} r' \xrightarrow{o/y}_{\mathcal{A}} r''}{(q, r) \xrightarrow{x/y}_{\alpha} (q', r'')} \quad \frac{\nexists i \in I : r \xrightarrow{i/x}_{\mathcal{A}}}{(q, r) \xrightarrow{x/\perp}_{\alpha} (q, r)}$$

To understand how the mapper is utilized during learning, we follow the execution of a single input of a query. The learner produces an abstract input  $x$ , which it sends to the mapper. By inversely following abstraction function  $\lambda_{\mathcal{A}}$ , the mapper converts this to a concrete input  $i$  and updates its state via transition  $r \xrightarrow{i/x}_{\mathcal{A}} r'$ . The concrete input  $i$  is passed on to the teacher, which responds with a concrete output  $o$  according to  $q \xrightarrow{i/o}_{\mathcal{M}} q'$ . This triggers the transition  $r' \xrightarrow{o/y}_{\mathcal{A}} r''$  in which the mapper generates the corresponding abstract output  $y$  and updates its state again. The abstract output is then returned to the learner.

We notice that the abstraction function is utilized invertedly when translating inputs. More precisely, the abstract input that the learner provides is an output for the mapper. The translation from abstract to concrete involves picking an arbitrary concrete value that corresponds with the given abstract value. It could be that multiple concrete values can be picked, in which case all values should lead to the same abstract behavior in order to learn a deterministic abstract model. It can also be that no values correspond to the input abstraction, in which case, by the second rule,  $\perp$  is returned to the learner, without consulting the teacher. We define the *abstraction component* implementing  $\alpha_{\mathcal{A}}$  as the transducer which follows from the mapper  $\mathcal{A}$ , but inverts the abstraction of inputs.

From the perspective of a learner, a teacher for  $\mathcal{M}$  and abstraction component implementing  $\alpha_{\mathcal{A}}$  together behave exactly like a teacher for  $\alpha_{\mathcal{A}}(\mathcal{M})$ . If  $\alpha_{\mathcal{A}}(\mathcal{M})$  is

deterministic, then the learner will eventually succeed in learning a deterministic machine  $\mathcal{H}$  satisfying  $\alpha_{\mathcal{A}}(\mathcal{M}) \approx \mathcal{H}$ . In [2], also a *concretization* operator  $\gamma_{\mathcal{A}}$  is defined. This operator is the adjoint of the abstraction operator: it turns any abstract machine  $\mathcal{H}$  with symbols in  $X$  and  $Y$  into a concrete machine with symbols in  $I$  and  $O$ . If  $\mathcal{H}$  is deterministic then  $\gamma_{\mathcal{A}}(\mathcal{H})$  is weakly deterministic.

As shown in [2],  $\alpha_{\mathcal{A}}(\mathcal{M}) \leq \mathcal{H}$  implies  $\mathcal{M} \leq \gamma_{\mathcal{A}}(\mathcal{H})$ . This tells us that when we apply mapper  $\mathcal{A}$  during learning of some “large” Mealy machine  $\mathcal{M}$ , even though we may not be able to learn the behavior of  $\mathcal{M}$  exactly, the concretization  $\gamma_{\mathcal{A}}(\mathcal{H})$  of the learned abstract model  $\mathcal{H}$  is an over-approximation of  $\mathcal{M}$ , that is,  $\mathcal{M} \leq \gamma_{\mathcal{A}}(\mathcal{H})$ . Similarly to the abstraction component, a *concretization component* for mapper  $\mathcal{A}$  implements  $\gamma_{\mathcal{A}}$ . This component is again fully defined by a mapper, but handles abstraction of outputs invertedly. During model checking, the composition of the abstract model  $\mathcal{H}$  and the concretization component for  $\mathcal{A}$  provides us with an over-approximation of  $\mathcal{M}$ .

**Framework for Mapper Definition.** In order to apply our abstraction approach, we need an abstraction and a concretization component for a given mapper  $\mathcal{A}$ . We could implement these components separately in an arbitrary programming language, but then they would have to remain consistent with  $\mathcal{A}$ . Moreover, ensuring that translation in one component inverts the corresponding translation in the other is non-trivial, and difficult to maintain, as changes in one would have to be applied invertedly in the other.

We used an alternative approach, in which we first define a mapper and then derive the abstraction and concretization components automatically. To this end, we built a language for defining a mapper in terms of (finite) registers, and functions to encode transitions and outputs. Our language supports case distinctions with programming-style if-else-statements, and requires that every branch leads to exactly one output and updates registers exactly once, such that the translations are complete. Except for the restrictions of finiteness and determinism, our language has the expressiveness of a simple programming language and should thus be usable to abstract (and concretize reversely) a wide range of systems and protocols. Listing 1 shows the example of a mapper for a simple login system. The mapper stores the first password received, and compares subsequent passwords to it. The abstract passwords used by the learner are  $\{true, false\}$ , denoting a correct or incorrect password, respectively. At the first attempt, *true* invertedly maps to any concrete password, and *false* maps to  $\perp$ . Later on *true* invertedly maps to the value picked the first time, while *false* maps to any other value. For TCP, we define multiple abstraction functions for inputs and outputs, in terms of multiple parameters per input or output.

To derive the components, we need to implement the inverse of the abstraction function, for both inputs and outputs. This can be achieved using a constraint solver or by picking random concrete values and checking if they are translated to the sought after abstraction, by translating them with the mapper in the forward direction. The latter approach may be hard, as the concrete domain is usually very large. In addition to finding a corresponding concrete

---

**Listing 1.** A simple example mapper for a login system, in a simplified syntax

---

```

integer stored := -1;
map ENTER(integer password → boolean correct)
  if (stored = -1 ∧ password ≥ 0) ∨ stored = password then
    correct := true
  else
    correct := false
  end if
end map
update
  if stored = -1 ∧ password ≥ 0 then
    stored := password
  else
    stored := stored
  end if
end update

```

▷ Every path explicitly assigns a value

---

value, another purpose of executing abstractions invertedly is to test the abstraction: different possible values should lead to the same abstract behavior, as the learner cannot handle non-determinism. A constraint solver usually picks values in a very structured and deterministic way, which does not test the abstraction well. Picking concrete values randomly and checking the corresponding abstract value allows more control over obtaining a good test coverage, but is in general less scalable.

## 3 Learning Setup

### 3.1 TCP as a System Under Learning

In TCP there are two interacting entities, the *server* and the *client*, which communicate over a network through packets, comprising a header and application data. On both sides there is an application, initiating and using the connection through *socket calls*. Each entity is learned separately and is a SUL in the learning context. This SUL thus takes packets or socket calls as inputs. It can output packets or *timeout*, in case the system does not respond with any packet. RFC 793 [38] and its extensions, most notably [9, 34], specify the protocol.

Packets are defined as tuples, comprising sequence and acknowledgement numbers, a payload and flags. By means of abstraction, we reduce the size of sequence and acknowledgement number spaces. Each socket call also defines an abstract and concrete input. Whereas packet configurations are the same for both client and server, socket calls differ. The server can *listen* for connections and *accept* them, whereas the client can actively *connect*. Both parties can *send* and *receive* data, or *close* an established connection (specifically, a half-duplex close [9, p. 88]). The server can additionally *close* its listening socket. Values returned by socket calls are not in the output alphabet to reduce setup complexity.

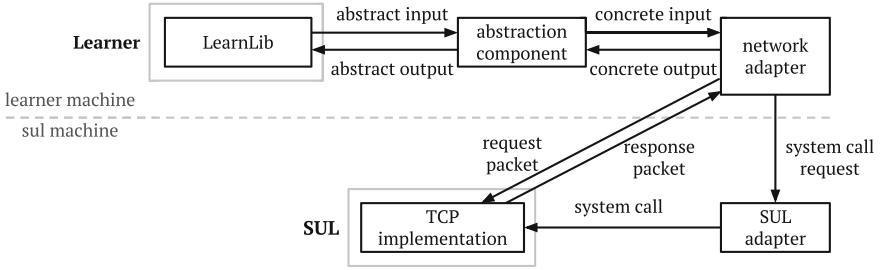


Fig. 1. Overview of the learning setup.

Figure 1 displays the learning setup used. The *learner* generates abstract inputs, representing packets or socket calls. The abstraction component concretizes each input by translating abstract parameters to concrete, and then updates its state. The concrete inputs are then passed on to the *network adapter*, which in turn transforms each input into a packet, sending it directly to the SUL, or into a socket call, which it issues to the *SUL adapter*. The *SUL adapter* runs on the same environment as the SUL and its sole role is to perform socket calls on the SUL. Each response packet generated by the SUL is received by the *network adapter*, which retrieves the concrete output from the packet or produces a *timeout* output, in case no packet was received within a predefined time interval. The output is then sent to the abstraction component, which computes the abstraction, updates its state again, and sends the abstract output to the learner.

The learner is based on LearnLib [39], a Java library implementing  $L^*$  based algorithms for learning Mealy machines. The abstraction component is also written in Java, and interprets and inverts a mapper. The network adapter is a Python program based on Scapy [41], Pcap [35], and Impacket [22]. It uses Scapy to craft TCP packets, and Scapy together with a Pcap and Impacket based sniffer to intercept responses. The network adapter is connected to the SUL adapter via a standard TCP connection. This connection is used for communicating socket calls to be made on the SUL. Finally, the SUL adapter is a program which performs socket calls on command, written in C to have low level access to socket options.

### 3.2 Viewing the SUL as a Mealy Machine

TCP implementations cannot be fully captured by Mealy Machines. To learn a model, we therefore need to apply some restrictions. As mentioned, the number of possible values for the sequence and acknowledgement numbers is reduced by means of abstractions. Furthermore, payload is limited to either 0 or 1 byte. Consequently, 1 byte of data is sent upon a *send*-call. Flags are limited to only the most interesting combinations, and we also abstract away from all other fields from the TCP layer or lower layers, allowing Scapy to pick default values.

TCP is also time-dependent. The SUL may, for instance, retransmit packets if they are not acknowledged within a specified time. The SUL may also reset if



it does not receive the acknowledgement after a number of such retransmissions, or if it remains in certain states for too long. The former we handled by having the *network adapter* ignore all retransmissions. For the latter, we verified that the learning queries were short enough so as not to cause these resets.

TCP is inherently concurrent, as a server can handle multiple connections at any time. This property is difficult to capture in Mealy Machines. Overcoming this, the SUL adapter ensures that at most one connection is accepted at any time by using a set of variables for locking and unlocking the *accept* and *connect*-calls. Only one blocking socket call can be pending at any time, but non-blocking socket calls can always be called.

Furthermore, the backlog size parameter defines the number of connections to be queued up for an eventual *accept*-call by the server SUL. The model grows linearly with the this parameter, while only exposing repetitive behavior. For this reason we set the backlog to the value 1.

### 3.3 Technical Challenges

We overcame several technical challenges in order to learn models. Resetting the SUL and setting a proper timeout value are solved similarly to [18].

Our tooling for sniffing packets sometimes missed packets generated by the SUL, reporting erroneous *timeout* outputs. This induced non-deterministic behavior, as a packet may or may not be caught, depending on timing. Each observation is therefore repeated three times to ensure consistency. Consistent outputs are cached to speed up learning, and to check consistency with new observations. It also allows to restart learning with reuse of previous observations.

In order to remove time-dependent behavior we use several TCP settings. Most notably, we disable slow acknowledgements and enable quick acknowledgements where possible (on Linux and FreeBSD). The intuition is that we want the SUL to send acknowledgements whenever they can be issued, instead of delaying them. We also had to disable syn cookies in FreeBSD, as this option caused generation of the initial sequence number in a seemingly time dependent way, instead of using fresh values. For Linux, packets generated by a *send*-call were occasionally merged with previous unacknowledged packets, so we could only learn a model by omitting *send*-call, although data packets could still be sent from the learner to the SUL.

### 3.4 Mapper Definition

The mapper is based on the work of Aarts et al. [5], and on the RFCs. Socket calls contain no parameters and do not need abstraction, so they are mapped simply with the identity relation. TCP packets are mapped by mapping their parameters individually. Flags are again retained by an identity relation. The sequence and acknowledgement numbers are mapped differently for inputs and outputs; input numbers are mapped to  $\{valid, invalid\}$ , and outputs are mapped to  $\{current, next, zero, fresh\}$ . After a connection is set up, the mapper keeps track of the sequence number which should be sent by the SUL and learner.

Valid inputs are picked according to this, whereas *current* and *next* represent repeated or incremented numbers, respectively. The abstract output *zero* simply denotes the concrete number zero, whereas *fresh* is used for all other numbers. If no connection is established, any sequence number is valid (as the RFCs then allow a fresh value), and the only valid acknowledgement number is zero.

Note that all concrete inputs with the same abstract value, should lead to an equivalent abstract behavior. Valid inputs are defined according to the RFC's. However, this is very hard for invalid inputs, as they may be accepted, ignored, they may lead to error recovery, or even undefined behavior. To learn the behavior for these inputs, abstractions should be defined precisely according to these behaviors, which is unfeasible to do by hand. As a result, we have excluded invalid inputs from the learning alphabet. To translate valid inputs, we first used a constraint solver which finds solutions for the transition relation. This is done by taking the disjunction of all path constraints, similar to symbolic execution techniques [25]. However, this did not test the abstraction well, as the constraint solver always picks zero if possible, for example. We therefore randomly picked concrete values to test with, and used values with the right corresponding abstract value, instead. Concrete values were picked with a higher probability if they were picked or observed previously during the same run, as well as the successors of those values. This approach sufficed to translate all values for our experiments.

## 4 Model Learning Results

Using the abstractions defined in Sect. 2, we learned models of the TCP client and server for Windows 8, Ubuntu 14.04 and FreeBSD 10.2. For testing we used the conformance testing algorithm described in [43] to generate efficient test suites which are parameterized by a middle section of length  $k$ . Generated exhaustively, these ensure learned model correctness, unless the respective implementation corresponds to a model with at least  $k$  more states. For each model, we first executed a random test suite with  $k$  of 4, up to 40000 tests for servers, and 20000 tests for clients. We then ran an exhaustive test suite with  $k$  of 2 for servers, respectively 3 for clients.

Table 1 describes the setting of each of these experiments together with statistics on learning and testing: (1) the number of states in the final model,

**Table 1.** Statistics for learning experiments

SUL		States	Hyp.	Memb. Queries	Tests to last Hyp.	Tests on last Hyp.
Client	Windows 8	13	2	1576	1322	50243
Server	Windows 8	38	10	11428	9549	65040
Client	Ubuntu 14.04	15	2	1974	15268	56174
Server	Ubuntu 14.04	57	14	17879	15681	66523
Client	FreeBSD 10.2	12	2	1456	1964	47387
Server	FreeBSD 10.2	55	18	22287	12084	75894

(2) the number of hypotheses found, (3) the total number of membership queries, (4) the total number of unique test queries run on the SUL before the last hypothesis, (5) the number of unique test queries run to validate the last hypothesis.

Figure 2 shows the model learned for the Windows 8 client. This model covers standard client behavior, namely connection setup, sending and receiving data and connection termination. Based on predefined access sequences, we identify each state with its analogous state in the RFC state diagram [38, p. 23], if such a state exists. Transitions taken during simulated communication between a Windows client and a server are colored green. These transitions were identified during model checking, on which we expand in Sect. 5.

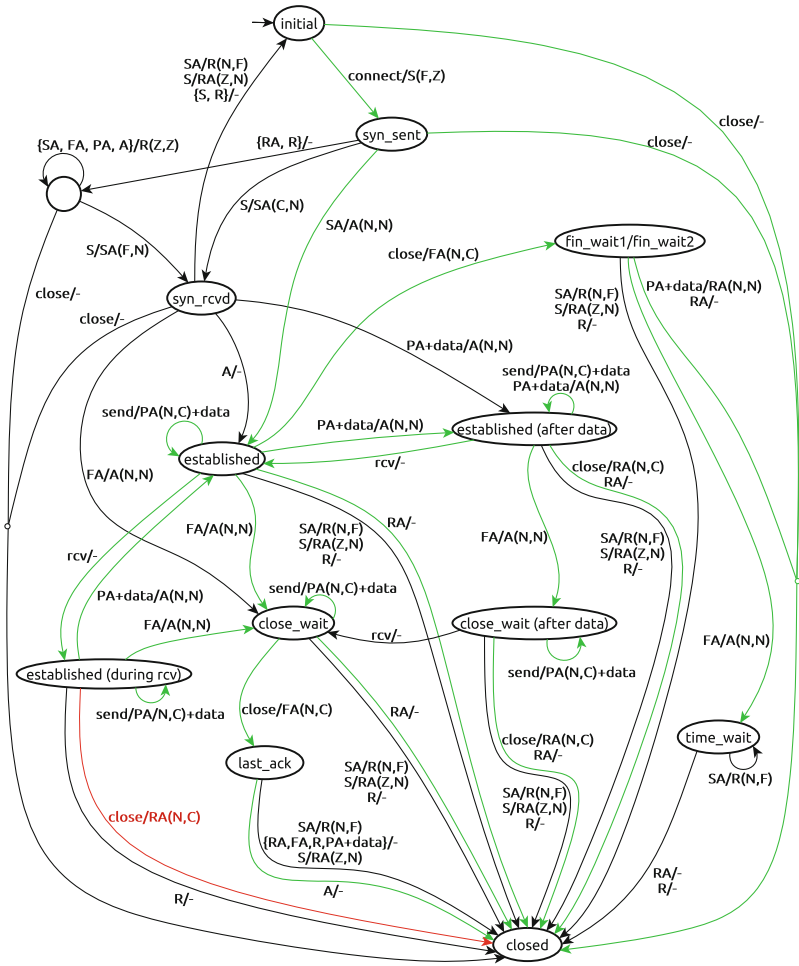
Table 1 shows that the models for the Linux and FreeBSD servers have more states than for Windows, and all models have more states than described in the specification. We attribute this to several factors. We have already mentioned that model sizes grow linearly with the value of the backlog-parameter. While we set it to 1, the setting is overridden by operating system imposed minimum value of 2 for FreeBSD and Linux. Moreover, SUL behavior depends on blocking system calls and on whether the receive buffer is empty or not. Although specified, this is not modelled explicitly in the specification state diagram. As an example, the ESTABLISHED and CLOSE WAIT states from the standard each have multiple corresponding states in the model in Fig. 2.

**Non-conformance of Implementations.** Inspection of learned models revealed several cases of non-conformance to RFC’s in the corresponding implementations.

A first non-conformance involves terminating an established connection with a CLOSE. The resulting output should contain a FIN if all data up to that point has been received. If there is data not yet received, the output should contain a RST, which would signal to the other side an aborted termination [9, p. 88]. Windows does not conform to this, as a CLOSE can generate a RST instead of a FIN even in cases where there is no data to be received, namely, in states where a RCV call is pending. Figure 2 marks this behavior in red. FreeBSD implementations are also non-compliant, as they always generate FIN packets on a CLOSE, regardless if all data has been received. This would arguably fall under the list of common bugs [34], namely “Failure to send a RST after Half Duplex Close”. The learned Linux models fully comply to these specifications.

A second non-conformance has to do with the processing of SYN packets. On receiving a SYN packet in a synchronized state, if the sequence number is in “the window” (as it always is, in our case), the connection should be reset (via a corresponding RST packet) [38, p. 71]. Linux implementations conform for SYN packets but not for SYN+ACK packets, to which they respond by generating an acknowledgement with no change of state. Both Windows and FreeBSD respect this specification.

We note a final non-conformance in Windows implementations. In case the connection does not exist (CLOSED), a reset should be sent in response to any



**Fig. 2.** Learned model for Windows 8 TCP Client. To reduce the size of the diagram, we eliminate all self loops with *timeout* outputs. We replace flags and abstractions by their capitalized initial letter, hence use *s* for *syn*, *a* for *ack*, *n* for *next*, etc. We omit input parameter abstractions, since they are the same for all packets, namely *valid* for both sequence and acknowledgement numbers. Finally, we group inputs that trigger a transition to the same state with the same output. Timeouts are denoted by ‘-’.

incoming packet except for another reset [38, p. 36], but Windows 8 sends nothing. FreeBSD can be configured to respond in a similar way to Windows, by changing the *blackhole setting*.<sup>1</sup> This behavior is claimed to provide “some degree of protection against stealth scans”, and is thus intentional.

<sup>1</sup> <https://www.freebsd.org/cgi/man.cgi?query=blackhole>.

## 5 Model Checking Results

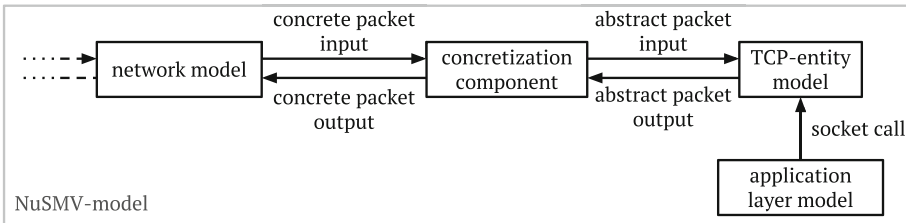
### 5.1 Model Checking the Learned Behaviour

We analyzed the learned models of TCP implementations using the model checker NuSMV [15]. We composed pairs of learned client and server models with a hand-made model of a non-lossy network, which simply delivers output from one entity as input for the other entity. Since the abstract input and output domains are different, the abstract models cannot communicate directly, and so we had to encode the concretized models within NuSMV code. We wrote a script that translated the abstract Mealy machine models from LearnLib to NuSMV modules, and another script that translated the corresponding mappers to NuSMV modules. TCP entities produce abstract outputs, which are translated to concrete. The network module then passes along such concrete messages. Before being delivered to the other entity, these messages are again transformed into abstract inputs. By encoding mapper functions as relations, NuSMV is able to compute both the abstraction function and its inverse, i.e., act as a concretization component. The global structure of the model is displayed in Fig. 3.

In Mealy machines, transitions are labeled by an input/output pair. In NuSMV transitions carry no labels, and we also had to split the Mealy machine transitions into a separate input and output part in order to enable synchronization with the network. Thus, a single transition  $q \xrightarrow{i/o} q'$  from a (concrete) Mealy machine is translated to a pair of transitions in NuSMV:

$$(loc = q, in = .., out = ..) \rightarrow (loc = q, in = i, out = ..) \rightarrow (loc = q', in = i, out = o).$$

Sequence and acknowledgement numbers in the implementations are 32-bit numbers, but were restricted to 3-bit numbers to reduce the state space. Whereas concrete messages are exchanged from one entity to the other, socket call inputs from the application are simulated by allowing system-calls to occur non-deterministically. A simplification we make is that we do not allow parallel actions: an action and all resulting packets have to be fully processed until another action can be generated. Consequently, there can be at most one packet in the composite model at any time. For example, once a three way handshake



**Fig. 3.** Schematic overview of NuSMV-model. Only half of the setup is shown in detail, as the model is symmetric and another TCP-entity model is connected to the network.

is initiated between a client and a listening server via a *connect*-call, no more system-calls can be performed until the handshake is finalized.

## 5.2 Checking Specifications

After a model is composed, the interaction between TCP entities can be analyzed using the NuSMV model checker. However, it is important to realize that, since we used abstractions, the learned models of TCP servers and clients are over-approximations of the actual behaviors of these components. If Mealy machine  $\mathcal{M}$  models the actual behavior of a component,  $\mathcal{A}$  is the mapper used, and  $\mathcal{H}$  is the abstract model that we learned then, as explained in Sect. 2, correctness of  $\mathcal{H}$  implies  $\mathcal{M} \leq \gamma_{\mathcal{A}}(\mathcal{H})$ . Since  $\gamma_{\mathcal{A}}(\mathcal{H})$  is weakly deterministic, in this case there exists a forward simulation relation from  $\mathcal{M}$  to  $\gamma_{\mathcal{A}}(\mathcal{H})$ . This forward simulation is preserved by the translation from Mealy machines to NuSMV. Results from Grumberg and Long [20] then imply that, for any  $\forall\text{CTL}^*$ -formula (which includes all LTL-formulas) we can transfer model checking results for  $\gamma_{\mathcal{A}}(\mathcal{H})$  to the (unknown) model  $\mathcal{M}$ . Since simulations are preserved by composition, this result even holds when  $\gamma_{\mathcal{A}}(\mathcal{H})$  is used as a component in a larger model.

Another essential point is that only a subset of the abstract inputs is used for learning. Hence invalid inputs (i.e. inputs with invalid parameters) are not included in our models. Traces with these inputs can therefore not be checked. Hence, the first property that we must check is a global invariant that asserts that invalid inputs will never occur. In case they do, NuSMV will provide a counterexample, which is used to find the cause of invalidity. During our initial experiments, NuSMV found several counterexamples showing that invalid inputs may occur. Based on analysis of these counterexamples we either refined/corrected the definition of one of the mappers, or we discovered a counterexample for the correctness of one of the abstract models. After a number of these iterations, we obtained a model in which invalid inputs can no longer occur. As mappers construction is done manually, these iterations are also not yet automated.

With only valid inputs, composite model may be checked for arbitrary  $\forall\text{CTL}^*$  formulas. Within these formulas, we may refer to input and output packets and their constituents (sequence numbers, acknowledgements, flags,...). This yields a powerful language for stating properties, illustrated by a few examples below. These formulas are directly based on the RFC's.

Many properties that are stated informally in the RFC's refer to control states of the protocol. These control states, however, cannot be directly observed in our black-box setting. Nevertheless, we can identify states, e.g. based on inputs and outputs leading to and from it. For example, we base the proposition *established* on RFC 793, which states that: "The connection becomes 'established' when sequence numbers have been synchronized in both directions" [38, p. 11], and that only a CLOSE or ABORT socket call or incoming packets with a RST or FIN can make an entity leave the ESTABLISHED state [38, Sect. 3.9].

We first show a simple safety formula checking desynchronization: if one entity is in the ESTABLISHED state, the other cannot be in SYN\_SENT and TIME\_WAIT:

$$G\neg(tcp1 - state = established \wedge (tcp2 - state = syn\_sent \vee tcp2 - state = time\_wait))$$

The next specification considers terminating an established connection with a CLOSE-input. The output should contain a FIN, except if there is unread data (in which case it should contain a RST). This corresponds to the first non-conformance case explained in Sect. 4. The specification is captured by the following formula, in which T is the triggered-operator as defined in NuSMV.

$$G(state = established \rightarrow ((input = rcv \text{ T } input \neq \text{ packet with data}) \wedge input = close) \rightarrow (F output = \text{ packet with FIN}))$$

We have formalized and checked, in a similar way, specifications for all other non-conforming cases as well as many other specifications.

We have also checked which transitions in the abstract models are reachable in the composed system. For every transition, we take its input and starting state, and check whether they can occur together. In this way we can find the reachable parts of model. This proves useful when analysing models, as the reachable parts likely harbor bugs with the most impact. Similarly, comparing reachable parts helps reveal the most relevant differences between implementations. The first and third non-conformances in Sect. 4 occur in the reachable parts of the respective models. Figure 2 marks these parts in green.

## 6 Conclusions and Future Work

We combined model learning, model checking and abstraction techniques to obtain and analyze models of Windows, Linux and FreeBSD TCP server and client implementations. Composing these models together with the model of a network allowed us to perform model checking over the composite setup and verify that any valid number generated by one TCP entity is seen as valid number by the other TCP entity. We have also identified breaches of the RFC's in all operating systems, and confirmed them by formulating temporal specification and checking them. Our work suggests several directions for future work.

Based on our understanding of TCP, we manually defined abstractions (mappers) that made it possible to learn models of TCP implementations. Getting the mapper definitions right turned out to be tricky. In fact, we had to restrict our learning experiments to *valid* abstractions of the sequence and acknowledgement numbers. This proved limiting when searching for interesting rules to model check, like for example those that would expose known implementation bugs. Such rules often concern invalid parameters, which do not appear in the models we learned. Additionally, we had to manually refine our mapper due to counterexamples found by the model checker. Learning algorithms that construct the abstractions automatically could potentially solve this problem. We hope that extensions of the learning algorithms for register automata as

implemented in the Tomte [4] and RALib [12] tools will be able to construct abstractions for TCP fully automatically.

Our work was severely restricted by the lack of expressivity of Mealy machines. In order to squeeze the TCP implementation into a Mealy machine, we had to eliminate timing based behavior as well as re-transmissions. Other frameworks for modeling state machines might facilitate modelling these aspects. Obviously, we would also need learning algorithms capable of generating such state machines. There has been some preliminary work on extending learning algorithms to timed automata [19, 47], and to I/O transition systems [3, 48], with the additional benefit of approximate learning. Approximate learning learns an upper and lower boundary to the behaviour of the system, instead of an exact model. This may allow to abstract away the corner-cases in the model and the mapper, if they are not relevant for the specifications. Extensions of this work could eliminate some of the restrictions that we encountered.

## References

1. Aarts, F.: Tomte: Bridging the Gap between Active Learning and Real-World Systems. Ph.D. thesis, Radboud University Nijmegen, October 2014
2. Aarts, F., Jonsson, B., Uijen, J., Vaandrager, F.W.: Generating models of infinite-state communication protocols using regular inference with abstraction. *Formal Methods Syst. Des.* **46**(1), 1–41 (2015)
3. Aarts, F., Vaandrager, F.: Learning I/O automata. In: Gastin, P., Laroussinie, F. (eds.) CONCUR 2010. LNCS, vol. 6269, pp. 71–85. Springer, Heidelberg (2010)
4. Aarts, F., Fiterau-Brostean, P., Kuppens, H., Vaandrager, F.: Learning register automata with fresh value generation. In: Leucker, M., Rueda, C., Valencia, F.D. (eds.) ICTAC 2015. LNCS, vol. 9399, pp. 165–183. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-25150-9\\_11](https://doi.org/10.1007/978-3-319-25150-9_11)
5. Aarts, F., Jonsson, B., Uijen, J., Vaandrager, F.W.: Generating models of infinite-state communication protocols using regular inference with abstraction. *Formal Methods Syst. Des.* **46**(1), 1–41 (2015)
6. Angluin, D.: Learning regular sets from queries and counterexamples. *Inf. Comput.* **75**(2), 87–106 (1987)
7. Baier, C., Katoen, J.-P.: Principles of Model Checking. MIT Press, Cambridge (2008)
8. Berendsen, J., Gebremichael, B., Vaandrager, F.W., Zhang, M.: Formal specification and analysis of Zeroconf using Uppaal. *ACM Trans. Embed. Comput. Syst.* **10**(3), 1–32 (2011)
9. Braden, R.: RFC 1122 Requirements for Internet Hosts - Communication Layers. Internet Engineering Task Force, October 1989
10. Brinksma, E., Mader, A.: On verification modelling of embedded systems. Technical report TR-CTIT-04-03, Centre for Telematics and Information Technology, Univ. of Twente, The Netherlands, January 2004
11. Bruns, G., Staskauskas, M.G.: Applying formal methods to a protocol standard and its implementations. In: Proceedings International Symposium on Software Engineering for Parallel and Distributed Systems (PDSE 1998), 20–21 April 1998, Kyoto, Japan, pp. 198–205. IEEE Computer Society (1998)



12. Cassel, S.: Learning Component Behavior from Tests: Theory and Algorithms for Automata with Data. Ph.D. thesis, University of Uppsala (2015)
13. Chalupar, G., Peherstorfer, S., Poll, E., de Ruitter, J.: Automated reverse engineering using Lego. In: Proceedings 8th USENIX Workshop on Offensive Technologies (WOOT 2014), San Diego, California, Los Alamitos, CA, USA, IEEE Computer Society, August 2014
14. Cho, C.Y., Babic, D., Shin, E.C.R., Song, D.: Inference and analysis of formal models of botnet command and control protocols. In: Al-Shaer, E., Keromytis, A.D., Shmatikov, V. (eds.) ACM Conference on Computer and Communications Security, pp. 426–439. ACM (2010)
15. Cimatti, A., Clarke, E., Giunchiglia, E., Giunchiglia, F., Pistore, M., Roveri, M., Sebastiani, R., Tacchella, A.: NuSMV 2: an opensource tool for symbolic model checking. In: Brinksma, E., Larsen, K.G. (eds.) CAV 2002. LNCS, vol. 2404, pp. 359–364. Springer, Heidelberg (2002)
16. Clarke, E.M., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement for symbolic model checking. *J. ACM* **50**(5), 752–794 (2003)
17. Clarke, E.M., Grumberg, O., Peled, D.: Model Checking. MIT Press, Cambridge (1999)
18. Fiterău-Broștean, P., Janssen, R., Vaandrager, F.: Learning fragments of the TCP network protocol. In: Lang, F., Flammini, F. (eds.) FMICS 2014. LNCS, vol. 8718, pp. 78–93. Springer, Heidelberg (2014)
19. Grinchtein, O., Jonsson, B., Leucker, M.: Learning of event-recording automata. *Theor. Comput. Sci.* **411**(47), 4029–4054 (2010)
20. Grumberg, O., Long, D.E.: Model checking and modular verification. *ACM Trans. Program. Lang. Syst.* **16**(3), 843–871 (1994)
21. Holzmann, G.J.: The SPIN Model Checker: Primer and Reference Manual. Addison Wesley, Reading (2004)
22. Impacket. <http://www.coresecurity.com/corelabs-research/open-source-tools/impacket>. Accessed 28 Jan 2016
23. Isberner, M: Foundations of Active Automata Learning: An Algorithmic Perspective. Ph.D. thesis, Technical University of Dortmund (2015)
24. Jhala, R., Majumdar, R.: Software model checking. *ACM Comput. Surv.* **41**(4), 21:1–21:54 (2009)
25. King, J.C.: Symbolic execution and program testing. *Commun. ACM* **19**(7), 385–394 (1976)
26. van Langevelde, I., Romijn, J.M.T., Goga, N.: Founding FireWire bridges through Promela prototyping. In: 8th International Workshop on Formal Methods for Parallel Programming: Theory and Applications (FMPPTA). IEEE Computer Society Press, April 2003
27. Lee, D., Yannakakis, M.: Principles and methods of testing finite state machines – a survey. *Proc. IEEE* **84**(8), 1090–1123 (1996)
28. Lockfeer, L., Williams, D.M., Fokkink, W.J.: Formal specification and verification of TCP extended with the window scale option. In: Lang, F., Flammini, F. (eds.) FMICS 2014. LNCS, vol. 8718, pp. 63–77. Springer, Heidelberg (2014)
29. Loiseaux, C., Graf, S., Sifakis, J., Boujjani, A., Bensalem, S.: Property preserving abstractions for the verification of concurrent systems. *Formal Methods Syst. Des.* **6**(1), 11–44 (1995)
30. Meinke, K., Sindhu, M.A.: Incremental learning-based testing for reactive systems. In: Gogolla, M., Wolff, B. (eds.) TAP 2011. LNCS, vol. 6706, pp. 134–151. Springer, Heidelberg (2011)

31. Musuvathi, M., Engler, D.R.: Model checking large network protocol implementations. In: Morris, R., Savage, S. (eds.) 1st Symposium on Networked Systems Design and Implementation (NSDI 2004), March 29–31, 2004, San Francisco, California, USA, Proceedings, pp. 155–168. USENIX (2004)
32. Musuvathi, M., Park, D.Y.W., Chou, A., Engler, D.R., Dill, D.L.: CMC: A pragmatic approach to model checking real code. In: Culler, D.E., Druschel, P. (eds.) 5th Symposium on Operating System Design and Implementation (OSDI 2002), Boston, Massachusetts, USA, December 9–11, 2002. USENIX Association (2002)
33. Nerode, A.: Linear automaton transformations. *Proc. Am. Math. Soc.* **9**(4), 541–544 (1958)
34. Paxson, V., Allman, M., Dawson, S., Fenner, W., Griner, J., Heavens, I., Lahey, K., Semke, J., Volz, B.: Known TCP Implementation Problems. RFC 2525 (Informational), March 1999
35. Pcap. <http://www.coresecurity.com/corelabs-research/open-source-tools/pcapy>. Accessed 28 Jan 2016
36. Peled, D., Vardi, M.Y., Yannakakis, M.: Black box checking. *J. Autom. Lang. Comb.* **7**(2), 225–246 (2002)
37. Popper, K.R.: *Logik der Forschung*. Julius Springer Verlag, Vienna (1935)
38. Postel, J.: Transmission control protocol. RFC 793 (Standard), Updated by RFCs 1122, 3168, September 1981
39. Raffelt, H., Steffen, B., Berg, T., Margaria, T.: LearnLib: a framework for extrapolating behavioral models. *STTT* **11**(5), 393–407 (2009)
40. de Ruiter, J., Poll, E.: Protocol state fuzzing of tls implementations. In: 24th USENIX Security Symposium (USENIX Security 15), pp. 193–206. SENIX Association, Washington, D.C., August 2015
41. Scapy. <http://www.secdev.org/projects/scapy/>. Accessed 28 Jan 2016
42. Shahbaz, M., Groz, R.: Analysis and testing of black-box component-based systems by inferring partial models. *Softw. Test. Verif. Reliab.* **24**(4), 253–288 (2014)
43. Smeenk, W., Moerman, J., Vaandrager, F., Jansen, D.N.: Applying automata learning to embedded control software. In: Butler, M., Conchon, S., Zaidi, F. (eds.) ICFEM 2015. LNCS, vol. 9407, pp. 67–83. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-25423-4\\_5](https://doi.org/10.1007/978-3-319-25423-4_5)
44. Steffen, B., Howar, F., Merten, M.: Introduction to active automata learning from a practical perspective. In: Bernardo, M., Issarny, V. (eds.) SFM 2011. LNCS, vol. 6659, pp. 256–296. Springer, Heidelberg (2011)
45. Stoelinga, M.: Fun with FireWire: A comparative study of formal verification methods applied to the IEEE 1394 root contention protocol. *Formal Aspects Comput. J.* **14**(3), 328–337 (2003)
46. Verleg, P.: Inferring SSH state machines using protocol state fuzzing. Master thesis, Radboud University (2016)
47. Verwer, S.: Efficient Identification of Timed Automata – Theory and Practice. Ph.D. thesis, Delft University of Technology, March 2010
48. Volpato, M., Tretmans, J.: Active learning of nondeterministic systems from an ioco perspective. In: Margaria, T., Steffen, B. (eds.) ISOla 2014, Part I. LNCS, vol. 8802, pp. 220–235. Springer, Heidelberg (2014)