Formal Methods in System Design manuscript No.
(will be inserted by the editor)

Wireless Protocol Validation Under Uncertainty

Jinghao Shi[1] · **Shuvendu K. Lahiri**[2] · **Ranveer Chandra**[2] · **Geoffrey Challen**[1]

the date of receipt and acceptance should be inserted later

Abstract Runtime validation of wireless protocol implementations cannot always employ direct instrumentation of the device under test (DUT). The DUT may not implement the required instrumentation, or the instrumentation may alter the DUT's behavior when enabled. Wireless sniffers can monitor the DUT's behavior without instrumentation, but they introduce new validation challenges. Losses caused by wireless propagation prevent sniffers from perfectly reconstructing the actual DUT packet trace. As a result, accurate validation requires distinguishing between specification deviations that represent implementation errors and those caused by sniffer uncertainty.

We present a new approach enabling sniffer-based validation of wireless protocol implementations. Beginning with the original protocol monitor state machine, we automatically and completely encode sniffer uncertainty by selectively adding non-deterministic transitions. We characterize the NP-completeness of the resulting decision problem and provide an exhaustive algorithm for searching over all mutated traces. We also present practical protocol-oblivious heuristics for searching over the most likely mutated traces. We have

This work was published, in part, in Runtime Verification (RV) 2016 [32].

Jinghao Shi
jinghaos@buffalo.edu

Shuvendu K. Lahiri
shuvendu@microsoft.com

Ranveer Chandra
ranveer@microsoft.com

Geoffrey Challen
challen@buffalo.edu

[1]University at Buffalo, Buffalo, NY 14260, USA
[2]Microsoft Research, Redmond, WA 98052, USA

implemented our framework and show that it can accurately identify implementation errors in the face of uncertainty.

Keywords Runtime Verification · Wireless Protocol · Sniffer · Uncertainty

1 Introduction

Custom wireless protocols are often designed and deployed to meet the specific performance and power needs of special-purpose wireless devices. Examples include Google Iris contact lenses [17], Xbox One wireless controllers [36], and Google Chromecast [35]. Validating that device implementations work correctly is critical to achieve the design goals of the wireless protocol and also prevent bugs in shipped products [9,15,11].

Runtime validation of the protocol implementations on such devices is challenging because collecting traces from the device under test (DUT) is often infeasible. The resource limitations of embedded or battery-powered devices may cause them to not provide trace collecting capabilities. DUT may contain proprietary hardware or firmware that hides the implementation details and prevents testers from collecting traces through source code instrumentation. Even when collecting trace directly from the DUT is possible, the overhead it causes may alter the behavior of the DUT due to the observer effect [28], threatening the validation results.

An attractive alternative is to use wireless sniffers to record traffic generated by the DUT during testing. Sniffers do not require direct access to the DUT or the need to alter its behavior. However, due to the fundamentally unpredictable nature of wireless communications, the packets captured by the sniffer will not exactly match those received by the DUT. The sniffer may miss packets that the DUT received, or receive packets that the DUT missed. This is true even when using multiple sniffers [8,25,3], a sniffer with multiple antennas [31], or in isolated wireless environments.

Since the sniffer trace may not perfectly match the actual trace, uncertainty arises during protocol implementation validation. For example, if the DUT fails to respond correctly to a packet in the sniffer trace, it may either because the DUT's implementation is incorrect, or the DUT did not actually receive the packet, or the DUT's response was missed by the sniffer. Whenever the DUT's behavior does not match the specification, there are now two potential explanations: either the DUT's implementation is wrong, or the sniffer trace is inaccurate. Accurate validation requires distinguishing between these two causes.

We present a new technique that enables validation of protocol implementations using wireless sniffers. Given a monitor state machine representing the protocol being validated, we describe a systematic transformation that adds non-deterministic transitions to incorporate uncertainty introduced by the sniffer. This augmented validation state machine implicitly defines a set of mutated traces, each satisfying the original state machine with a specific likelihood. If the set of mutated traces is empty, the implementation definitely

violates the protocol. Searching over all the mutated traces is NP-complete, but the approach can be made practical by applying protocol-oblivious heuristics that limit the search to likely mutated traces.

Our paper makes the following contributions:

1. To the best of our knowledge, we are the first to identify the uncertainty problem caused by sniffers in validating wireless protocol implementations.
2. We formalize the problem using a nondeterministic state machine that systematically and completely encodes the uncertainty of the sniffer trace.
3. We characterize the NP-completeness of the validation problem, and present two protocol-oblivious heuristics to prune the search space and make validation possible in practice.
4. We implement the validation framework and evaluate it using the NS-3 network simulator [29]. Our framework accurately identifies both synthetic and previously unknown violations in NS-3's implementations of the 802.11 and ARF protocols. We also applied our framework to a commercial product under development and was able to found three latent bugs that were not observed previously.

This paper is an extension of our work that appeared in RV 2016 [32]. We added the proof for Lemma 1, 2 and Theorem 1. We included new results that show the searching cost (Fig. 6), and details of the ARF protocol violations (Section 5.2). Further, we reported the application of our framework to a commercial product under development in Section 5.3.

2 Background and Motivating Example

We encountered the uncertainty problem while testing the protocol implementation of a popular wireless game controller. A custom wireless protocol was designed to meet the low latency and low power consumption goals. As is common industry practice, the protocol specification was then handed over to wireless chipset vendors for implementation. However, neither implementation details nor trace collection capabilities are provided in the shipped firmware due to intellectual property constraints and device resource limitation. Hence using sniffers to validate the protocol implementation is the only option.

We initially developed a tool to validate certain protocol properties over the sniffer trace, yet often found unacceptable amount of false alarms due to the incompleteness of the sniffer traces, making the tool virtually useless. It was clear that we needed to account for sniffer uncertainty.

To better understand the incompleteness of sniffer trace, consider the IEEE 802.11 (also known as Wi-Fi) transmitter (DUT) state machine shown in Fig. 1. After the DUT sends $DATA_i$—a data packet with sequence number i ($s_0 \rightarrow s_1$), it starts a timer and waits for the acknowledgment packet—Ack. The DUT either receives Ack within time T_o ($s_1 \rightarrow s_0$), or it sends $DATA'_i$—retransmission of $DATA_i$ ($s_1 \rightarrow s_2$). Similarly, the DUT either receives the

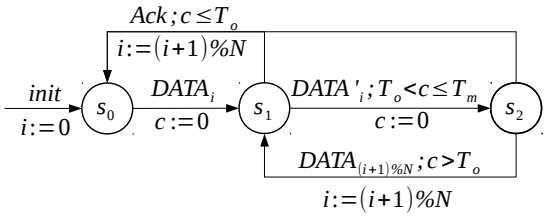

Fig. 1: **Monitor State Machine for 802.11 Transmitter.**

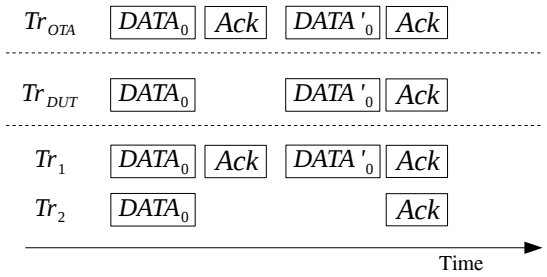

Fig. 2: **Uncertainty of Sniffer Observations.** Tr_{OTA} is the chronological sequence of packets sent by the DUT and the receiver. Tr_{DUT} is DUT's internal events. Tr_1 and Tr_2 are two examples of many possible sniffer traces.

Ack within T_o ($s_2 \rightarrow s_0$) or aborts transmission and moves on to next packet[1] ($s_2 \rightarrow s_1$).

Given a complete log of DUT's packet transmission and reception events, it is trivial to feed such a log into the state machine in Fig. 1 and validate the correctness of DUT's protocol implementation. However, due to DUT limitations we have described earlier, this complete log is not available. As a result, we seek to validate the DUT implementation using sniffers.

There are two fundamental properties in wireless communication that bring uncertainty to sniffer's observation: packet loss and physical diversity. The sniffer could either miss packets sent from or to the DUT due to packet loss, or overhear packets that are sent to but missed by the DUT due to physical diversity. Note that packet alternation needs not be considered due to packet level checksum mechanisms.

Consider a correct packet exchange sequence shown in Fig. 2. The DUT first sends $DATA_0$. Suppose the receiver receives $DATA_0$ and sends the *Ack* which the DUT does not receive. Eventually the DUT's timer fires and it sends $DATA'_0$. This time the $DATA'_0$ reaches receiver and the DUT also receives the *Ack*.

Now consider two possible traces that could have been overheard by a sniffer shown in Fig. 2. In first sniffer trace Tr_1 where the sniffer *overhears*

[1] To represent the state machine succinctly, our example assumes that the DUT retries at most once.

the first *Ack* packet, a validation *uncertainty* arises when the sniffer sees the $DATA'_0$: was the previous *Ack* missed by the DUT or is there a bug in DUT which causes it to retransmit even after receiving the *Ack*?

Similarly, consider the second possible sniffer trace Tr_2 where both the $DATA'_0$ and *Ack* packets were missed by the sniffer. During this period of time, it appears the DUT neither receives *Ack* for $DATA_0$ nor sends $DATA'_0$. Again, without any additional information it is impossible to disambiguate between the sniffer missing certain packets and a bug in DUT's retransmission logic.

Informally, the question we set out to answer in this paper is: given the protocol monitor state machine and the sniffer's observation with inherent uncertainty, how to accurately validate that the DUT behaves as specified?

3 Prerequisites and Problem Statement

3.1 Packet, Trace and Monitor State Machine

The alphabet of the monitor state machine is the finite set of all valid packets defined by the protocol, denoted as $\mathbb{P}$. A packet is a binary string of a finite number of bits, encoding interesting protocol attributes such as `src`, `dest`, `type`, `flags`, and physical layer information, such as `channel`, `modulation`, etc. The input of the state machine then corresponds to a time-ordered sequence of packets.

Definition 1 (Packet Trace) A *packet trace* is a finite sequence of $(timestamp, packet)$ tuple: $[(t_1, p_1), (t_2, p_2), \ldots, (t_n, p_n)]$ where $t_i \in \mathbb{Z}^+$ is the *discrete* timestamp and p_i is the packet observed at time t_i. The timestamps are strictly monotonically increasing: $t_i < t_{i+1}$ for $1 \leq i < n$.

In addition to timestamp monotonicity, we also require that adjacent packets do not overlap in time, $t_{i+1} - t_i > \texttt{airtime}(p_i)$ for $1 \leq i < n$, where `airtime()` calculates the time taken to transmit a packet. The timestamp represents the observer's local clock ticks, and need not to be synchronized among devices.

We use *timed automata* [1] to model the expected behaviors of the DUT. A timed automata is a finite state machine with timing constraints on the transitions: each transition can optionally start one or more timers, which can later be used to assert certain events should be seen before or after the time out event. We refer the readers to [1] for more details about timed automata.

Definition 2 (Monitor) A protocol monitor state machine S is a 7-tuple $\{\Sigma, \mathbb{S}, \mathbb{X}, s_0, C, E, G\}$, where:

- $\Sigma = \mathbb{P}$ is the finite input alphabet.
- $\mathbb{S}$ is a non-empty, finite set of states. $s_0 \in \mathbb{S}$ is the initial state.
- $\mathbb{X}$ is the set of boolean variables. We use $v = \{x \leftarrow true/false \mid x \in \mathbb{X}\}$ to denote an assignment of the variables. Let $\mathbb{V}$ be the set of such values v.

- C is the set of clock variables. A *clock variable* can be reset along any state transitions. At any instant, reading a clock variable returns the time elapsed since last time it was reset.
- G is the set of guard conditions defined inductively by

$$g := true \mid c \leq T \mid c \geq T \mid x \mid \neg g \mid g_1 \wedge g_2$$

  where $c \in C$ is a clock variable, T is a constant, and x is a variable in $\mathbb{X}$. A transition can choose not to use guard conditions by setting g to be *true*.
- $E \subseteq \mathbb{S} \times \mathbb{V} \times \mathbb{S} \times \mathbb{V} \times \Sigma \times G \times \mathscr{P}(C)$ gives the set of transitions. $\langle s_i, v_i, s_j, v_j, p, g, C' \rangle \in E$ represents that if the monitor is in state s_i with variable assignments v_i, given the input tuple (t, p) such that the guard g is satisfied, the monitor can transition to a state s_j with variable assignments v_j, and reset the clocks in C' to 0.

A tuple (t_i, p_i) in the packet trace means the packet p_i is presented to the state machine at time t_i. As an example, the monitor state machine illustrated in Fig. 1 can be formally defined as follows:

- $\Sigma = \{DATA_i, DATA'_i, Ack \mid 0 \leq i < N\}$.
- Clock variables $C = \{c\}$. The only clock variable c is used for acknowledgment time out.
- $\mathbb{X} = \{i\}$, as a variable with $log(N) + 1$ bits to count from 0 to N.
- Guard constraints $G = \{c \leq T_o, c > T_o, T_o < c \leq T_m\}$. T_o is the acknowledgment time out value, and $T_m > T_o$ is the maximum delay allowed before the retransmission packet gets sent. In this particular case, T_o can be arbitrary large but not infinity in order to check the liveness of the DUT.

Definition 3 (Trace Rejection and Acceptance) A monitor state machine S *rejects a trace* Tr if there exists a prefix of Tr such that all states reachable after consuming the prefix have no valid transitions for the next (t, p) input. Otherwise, S *accepts* Tr (or Tr *satisfies* S).

The monitor state machine defines a *timed language* L which consists of all valid packet traces that can be observed by the DUT. We now give the definition of protocol *compliance* and *violation*.

Definition 4 (Violation and Compliance) Suppose $\mathbb{T}$ is the set of all possible packet traces collected from DUT, and S is the state machine specified by the protocol. The DUT *violates* the protocol specification if there exists a packet trace $Tr \in \mathbb{T}$ such that S rejects Tr. Otherwise, the DUT is *compliant* with the specification.

The focus of this paper is to determine whether a *given* sniffer trace Tr is evidence of a violation.

3.2 Mutation Trace

As shown in the motivation example in Fig. 2, a sniffer trace may either miss packets that are present in DUT trace, or contain extra packets that are missing in DUT trace. Note that in the latter case, those extra packets must be all sent *to* the DUT. This is because it is impossible for the sniffer to overhear packets sent from the DUT that were not actually sent by the DUT.

We formally capture this relationship with the definition of mutation trace.

Definition 5 (Mutation Trace) A packet trace Tr' is a mutation of sniffer trace Tr w.r.t a DUT if for all $(t, p) \in Tr \setminus Tr'$, $p.dest = DUT$, where $p.dest$ is the destination of packet p.

By definition, either $Tr' \supseteq Tr$ (hence $Tr \setminus Tr' = \emptyset$), or those extra packets in Tr but not in Tr' are all sent to the DUT. Note that Tr' may contain extra packets that are either sent to or received by the DUT.

A mutation trace Tr' represents a *guess* of the corresponding DUT packet trace given sniffer trace Tr. In fact, the DUT packet trace must be one of the mutation traces of the sniffer trace Tr.

Lemma 1 *Let Tr_{DUT} and Tr be the DUT and sniffer packet trace captured during the same protocol operation session, and $\mathcal{M}(Tr)$ be the set of mutation traces of Tr with respect to DUT, then $Tr_{DUT} \in \mathcal{M}(Tr)$.*

Proof Let $\Delta = Tr \setminus Tr_{DUT}$ be the set of packets that are in Tr but not in Tr_{DUT}. Recall that it is not possible for the sniffer to observe packets sent from the DUT that the DUT did not send. Therefore, all packets in Δ are sent *to* the DUT. That is, for all $(t, p) \in \Delta$, $p.dest = DUT$. By Definition 5, $Tr_{DUT} \in \mathcal{M}(Tr)$. □

3.3 Problem Statement

Lemma 1 shows that $\mathcal{M}(Tr)$ is a *complete* set of guesses of the DUT packet trace. Therefore, the problem of validating DUT implementation given a sniffer trace can be formally defined as follows:

Problem 1 VALIDATION
instance A protocol monitor state machine S and a sniffer trace Tr.
question Does there exist a mutation trace Tr' of Tr that satisfies S?

If the answer is no, a definite violation of the DUT implementation can be claimed. Nevertheless, if the answer is yes, S may still reject Tr_{DUT}. In other words, the conclusion of the validation can either be *definitely wrong* or *probably correct*, but not *definitely correct*. This is the fundamental limitation caused by the uncertainty of sniffer traces.

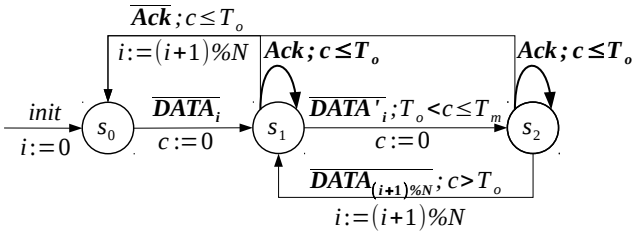

Fig. 3: **Augmented Monitor State Machine.** Augmented transitions are highlighted in bold face. $\overline{Pkt}$ means either ϵ or Pkt.

4 Validation Framework

4.1 Augmented State Machine

To deal with the inherent uncertainty of sniffer traces, we propose to systematically augment the original monitor state machine with non-deterministic transitions to account for the difference between the sniffer and DUT traces.

Before formally defining the augmented state machine, we first use an example to illustrate the basic idea. Fig. 3 shows the augmented state machine for 802.11 transmitter state machine shown in Fig. 1. For each existing transition (e.g., $s_0 \rightarrow s_1$), we add an *empty transition* with same clock guards and resetting clocks. This accounts for the possibility when such packet was observed by the DUT but missed by the sniffer. Additionally, for each transition triggered by a *receiving* packet (i.e., $p.dest = DUT$), such as $s_1 \rightarrow s_0$ and $s_2 \rightarrow s_0$, we add a *self transition* with the same trigger packet and clock guards, but an empty set of resetting clocks and no assignments to variables. This allows the state machine to make progress when the sniffer missed such packets.

There are two things to note. First, self transitions are added only for packets sent *to* the DUT, since the sniffer will not overhear packets *from* the DUT if they were not sent by the DUT. Second, no augmented transitions are added for the packets that are sent to DUT yet are missed by both the DUT and the sniffer, since such packets do not cause difference between the DUT and sniffer traces.

The augmented state machine in Fig. 3 will accept the sniffer packet traces Tr_1 and Tr_2 shown in Fig. 2. For instance, one accepting transition sequence on sniffer trace Tr_1 is $s_0 \rightarrow s_1 \rightarrow_s s_1 \rightarrow s_2 \rightarrow s_0$, and the sequence for Tr_2 is $s_0 \rightarrow s_1 \rightarrow_e s_2 \rightarrow s_0$, where $\rightarrow$ is the transition from the original state machine, $\rightarrow_e$ and $\rightarrow_s$ are the augmented empty and self transitions respectively.

We now formally define the augmented state machine.

Definition 6 (Augmented Monitor) An augmented state machine S^+ for a monitor state machine S is a 7-tuple $\{\boldsymbol{\Sigma}^+, \mathbb{S}, \mathbb{X}, s_0, C, \boldsymbol{E}^+, G\}$, where $\mathbb{S}, \mathbb{X}, s_0, C, G$ are the same as S. $\Sigma^+ = \{\epsilon\} \cup \Sigma$ is the augmented input alphabet with the empty symbol, and $E^+ \supset E$ is the set of transitions, which includes:

Algorithm 1 Obtain Augmented Transitions E^+ from E

1: **function** AUGMENT(E)
2: $E^+ := \emptyset$
3: **for all** $\langle s_i, v_i, s_j, v_j, p, g, C' \rangle \in E$ **do**
4: $E^+ := E^+ \cup \{\langle s_i, v_i, s_j, v_j, p, g, C' \rangle\}$ ▷ Type-0
5: $E^+ := E^+ \cup \{\langle s_i, v_i, s_j, v_j, \boldsymbol{\epsilon}, g, C' \rangle\}$ ▷ Type-1
6: **if** $p.dest = DUT$ **then**
7: $E^+ := E^+ \cup \{\langle s_i, v_i, \boldsymbol{s_i}, \boldsymbol{v_i}, p, g, \emptyset \rangle\}$ ▷ Type-2
8: **return** E^+

– E: existing transitions (**Type-0**) in S.
– E_1^+: empty transitions (**Type-1**) for transitions in E.
– E_2^+: self transitions **(Type-2)** for transitions triggered by receiving packets.

Algorithm 1 describes the process of transforming E into E^+. In particular, Line 4 adds existing transitions in E to E^+, while line 5 and 7 add Type-1 and Type-2 transitions to E^+ respectively. We have highlighted the elements of the tuple that differ from the underlying Type-0 transition. Note that in Type-2 transitions, both the state and the variables stay the same after the transition.

With augmented state machine S^+, we can use Type-1 transitions to non-deterministically infer packets missed by the sniffer, and use Type-2 transitions to consume extra packets captured by the sniffer but missed by the DUT.

An accepting run of S^+ on sniffer trace Tr yields a mutation trace Tr' which represents one possibility of the DUT trace. Specifically, Tr' can be obtained by adding missing packets indicated by Type-1 transitions to Tr, and removing extra packets indicated by Type-2 transitions from Tr

We show that the VALIDATION problem is equivalent to the satisfiability problem of Tr on S^+.

Theorem 1 *Let Tr be a sniffer trace, $\mathcal{M}(Tr)$ be Tr's mutation traces, S and S^+ be the original and augmented monitor state machine respectively. There exists a mutation trace $Tr' \in \mathcal{M}(Tr)$ that satisfies S if and only if Tr satisfies S^+.*

Proof Assume Tr satisfies S^+, and P is a sequence of accepting transitions, we construct a mutation trace Tr' using P and show that Tr' satisfies S.

Initially, let $Tr' = Tr$, then for each *augmented* transition $\langle s_i, v_i, s_j, v_j, \sigma, g, C' \rangle \in P$:

– If this is a Type-1 transition, add (t, p) to Tr', where t is a timestamp that satisfies g and p is the missing packet.
– If this is a Type-2 transition, remove corresponding (t, p) from Tr'.

It is obvious that Tr' is a mutation trace of Tr, since only receiving packets are removed in the process.

Now we show that there exists an accepting transition sequence P' of S^+ on input Tr' that does not contain augmented transitions. In particular, P' can be obtained by substituting all Type-1 transitions with corresponding original transitions, and removing all Type-2 transitions. Since P' does not contain augmented transitions, it is also an accepting transition sequence of S on input Tr', hence Tr' satisfies S.

On the other hand, assume $Tr' \in \mathcal{M}(Tr)$ and Tr' satisfies S. Suppose P' is the accepting transition sequences of S on input Tr'. We first note that P' is also the accepting transitions of S^+ on input Tr', since $E \subset E^+$.

We construct an accepting transition sequence P of S^+ on input Tr as follows.

- For each packet $p \in Tr' \setminus Tr$, substitute the transition $\langle s_i, v_i, s_j, v_j, p, g, C' \rangle$ with the corresponding Type-1 transition $\langle s_i, v_i, s_j, v_j, \epsilon, g, C' \rangle$.
- For each transition $\langle s_i, v_i, s_j, v_j, \sigma, g, C' \rangle$ followed by packet $p \in Tr \setminus Tr'$, add a Type-2 self transition $\langle s_j, v_j, s_j, v_j, p, g, \emptyset \rangle$. This is possible since Tr' is a mutation trace of Tr, thus for all $p \in Tr' \setminus Tr$, $p.src \neq DUT$.

Therefore, Tr satisfies S^+. □

By Theorem 1, the inherent uncertainty of the sniffer traces is explicitly represented by the augmented transitions, and can be systematically explored using the well established theory of state machine.

4.2 Problem Hardness

In this section, we show that the VALIDATION problem is NP-complete. In fact, the problem is still NP-complete even with only one type of augmented transitions.

Recall that Type-1 transitions are added because the sniffer may miss packets. Suppose an imaginary sniffer that is able to capture *every* packet ever transmitted, then only Type-2 transitions are needed since the sniffer may still overhear packets sent to the DUT. Similarly, suppose another special sniffer that would not overhear any packets sent to the DUT, then only Type-1 transitions are needed to infer missing packets.

We refer the augmented state machine that only has Type-0 and Type-1 transitions as S_1^+, and the augmented state machine that only has Type-0 and Type-2 transitions as S_2^+. And we show that each subproblem of determining trace satisfiability is NP-complete.

Problem 2 VALIDATION-1
Given that $Tr \setminus Tr_{DUT} = \emptyset$ (sniffer does not overhear packets).
instance Checker state machine S and sniffer trace Tr.
question Does S_1^+ accept Tr?

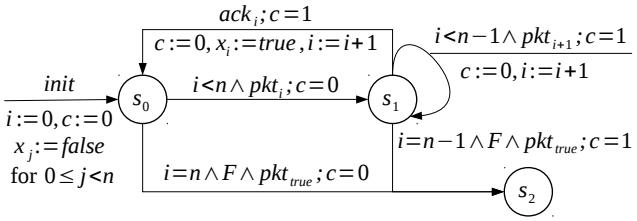

Fig. 4: **Monitor State Machine for SAT Problem.**

Problem 3 VALIDATION-2
Given that $Tr_{DUT} \subseteq Tr$ (sniffer does not miss packets).
instance Checker state machine S and sniffer trace Tr.
question Does S_2^+ accept Tr?

Lemma 2 *Both VALIDATION-1 and VALIDATION-2 are NP-complete.*

Proof First, note that the length of mutation trace Tr' is polynomial to the length of Tr because of the discrete time stamp and non-overlapping packets assumption. Therefore, given a state transition sequence as witness, it can be verified in polynomial time whether or not it is an accepting transition sequence, hence both VALIDATION-1 and VALIDATION-2 are in NP.

Next, we show how the SAT problem can be reduced to either one of the two problems. Consider an instance of SAT problem of a propositional formula F with n variables $x_0, x_1, \ldots, x_{n-1}$, we construct a corresponding protocol and its monitor state machine as follows.

The protocol involves two devices: the DUT (transmitter) and the endpoint (receiver). The DUT shall send a series of packets, $pkt_0, pkt_1, \ldots, pkt_{n-1}$. For each pkt_i, if the DUT receives the acknowledgment packet ack_i from the endpoint, it sets boolean variable x_i to be `true`. Otherwise x_i remains to be `false`. After n rounds, the DUT evaluate the formula F using the assignments and sends a special packet, pkt_{true}, if F is `true`. One round of the protocol operation can be completed in polynomial time since any witness of F can be evaluated in polynomial time.

The protocol monitor state machine S is shown in Fig. 4. Initially, all x_i is set to `false`. At state s_0, the DUT shall transmit pkt_i within a unit time, transit to s_1 and reset the clock along the transition. At state s_1, either the DUT receives the ack_i packet and set x_i to be `true` ($s_1 \rightarrow s_0$), or the DUT continues to transmit the next packet pkt_{i+1}. After n rounds, the state machine is s_0 or s_1 depending on whether ack_{n-1} is received by the DUT. In either case, the DUT shall evaluate F and transmit pkt_{true} if F is `true`.

Consider a sniffer trace $Tr_1 = \{(0, pkt_0), (2, pkt_1), (4, pkt_2), \ldots, (2(n-1), pkt_{n-1}), (2n, pkt_{true})\}$. That is, the sniffer only captures all pkt_i plus the final pkt_{true}, but none of ack_i. It is easy to see that F is satisfiable if S_1^+ accepts Tr_1. In particular, a successful run of S_1^+ on Tr_1 would have to guess, for each pkt_i, whether the Type-1 empty transitions should be used to infer

the missing ack_i packet, such that F is *true* at the end. Note that for Tr_1, no Type-2 self transitions can be used since all packets in Tr_1 are sent from the DUT. Therefore, the SAT problem of F can be reduced to the VALIDATION-1 problem of S_1^+ on sniffer trace Tr_1.

On the other hand, consider another sniffer trace $Tr_2 = \{(0, pkt_0), (1, ack_0), (2, pkt_1), (3, ack_1), \ldots, (2n-2, pkt_{n-1}), (2n-1, ack_{n-1}), (2n, pkt_{true}\}$. That is, the sniffer captures all n pair of pkt_i and ack_i packets and the final pkt_{true} packet. Similar to Tr_1, F is satisfiable if S_2^+ accepts Tr_2. A successful transition sequence of S_2^+ on Tr_2 must decide for each ack_i packet, whether Type-2 self transitions should be used, so that F can be evaluated as true at the end. Therefore, the SAT problem of F can also be reduced to the VALIDATION-2 problem of S_2^+ on sniffer trace Tr_2.

Since SAT is known to be NP-complete, both the VALIDATION-1 and the VALIDATION-2 problem are also NP-complete. □

The hardness statement of the general VALIDATION problem naturally follows Lemma 2.

Theorem 2 *VALIDATION is NP-complete.*

4.3 Searching Strategies

In this section, we present an *exhaustive* search algorithm of the accepting transition sequence of S^+ on sniffer trace Tr. It is guaranteed to yield an accepting sequence if there exists one, thus is exhaustive. In the next sections, we present heuristics to limit the search to accepting sequences of S^+ that require relatively fewer transitions from $E_1^+ \cup E_2^+$. Due to the NP-completeness of the problem, this also makes the algorithm meaningful in practice.

The main routines of the algorithm are shown in Algorithm 2. In the top level `SEARCH` routine, we first obtain the augmented state machine S^+, and then call the recursive `EXTEND` function with an empty prefix, the sniffer trace, and the S^+'s initial state. In the `EXTEND` function, we try to consume the first packet in the remaining trace using either Type-0, Type-1 or Type-2 transition. Note that we always try to use Type-0 transitions before other two augmented transitions (line 6). This ensures the first found mutation trace will have the most number of Type-0 transitions among all possible mutation traces. Intuitively, this means the search algorithm tries to utilize the sniffer's observation as much as possible before considering the packet loss by the sniffer or DUT.

Each of the extend functions either returns the mutation trace Tr', or *nil* if the search fails. In both `EXTEND-0` and `EXTEND-2` function, if there is a valid transition, we try to consume the next packet either by appending it to the prefix (line 13) or dropping it (line 26). While in `EXTEND-1`, we guess a missing packet without consuming the next real packet (line 20). Note that since only Type-0 and Type-2 consume packets, the recursion terminates if there is a valid Type-0 or Type-2 transition for the last packet (line 12 and line 25).

Algorithm 2 Exhaustive search algorithm of S^+ on Tr.

```
 1: function SEARCH(S, Tr)
 2:     S+ := AUGMENT(S)
 3:     return EXTEND([], Tr, S+.s0)
 4: function EXTEND(prefix, p::suffix, s)
 5:     if not LIKELY(prefix) then return nil[a]
 6:     for i ∈ [0, 1, 2] do
 7:         mutation := EXTEND-i(prefix, p::suffix, s)
 8:         if mutation ≠ nil then return mutation
 9:     return nil
10: function EXTEND-0(prefix, p::suffix, s)
11:     for ⟨s, s′, p⟩[b] ∈ E do
12:         if suffix = nil then return prefix@p
13:         mutation := EXTEND(prefix@p, suffix, s′)
14:         if mutation ≠ nil then return mutation
15:     return nil
16: function EXTEND-1(prefix, p::suffix, s)
17:     for all ⟨s, s′, q⟩ ∈ E1+ do
18:         if q.time > p.time then
19:             continue
20:         mutation := EXTEND(prefix@q, p::suffix, s′)
21:         if mutation ≠ nil then return mutation
22:     return nil
23: function EXTEND-2(prefix, p::suffix, s)
24:     for all ⟨s, s, p⟩ ∈ E2+ do
25:         if suffix = nil then return prefix
26:         mutation := EXTEND(prefix, suffix, s)
27:         if mutation ≠ nil then return mutation
28:     return nil
```

[a] This check should be ignored in the exhaustive algorithm.
[b] $\langle s, s', p\rangle$ is short for $\langle s, *, s', *, p, *, *\rangle$

It is not hard to see that Algorithm 2 terminates on any sniffer traces. Each node in the transition tree only has finite number of possible next steps, and the depth of Type-1 transitions is limited by the time available before the next packet (line 18).

4.4 Pruning Heuristics

In the face of uncertainty between a possible protocol violation and sniffer imperfection, augmented transitions provide the ability to blame the latter. The exhaustive nature of Algorithm 2 means that it always tries to blame sniffer imperfection whenever possible, making it reluctant to report true violations.

Inspired by the *directed model checking* [12] technique which is to mitigate the state explosion problem, we propose to enforce extra constraints on the mutation trace to restrict the search to only mutation traces with high

likelihood. The modified `EXTEND` function checks certain likelihood constraints on the prefix of the mutation trace before continuing (line 5), and stops the current search branch immediately if the prefix seems *unlikely*. Because of the recursive nature of the algorithm, other branches which may have a higher likelihood can then be explored.

The strictness of the likelihood constraint represents a trade-off between precision and recall of validation. The more strict the constraints are, the more false positive violations will potentially be reported, hence the lower the precision yet higher recall. On the contrary, the more tractable the constraints are, the more tolerant the search is to sniffer imperfection, hence the more likely that it will report true violations, thus higher precision but lower recall.

The exact forms of the constraints may depend on many factors, such as the nature of the protocol, properties of the sniffer, or domain knowledge. Next, we propose two *protocol oblivious* heuristics based on the sniffer loss probabilities and general protocol operations. Both heuristic contains parameters that can be fine tuned in practice.

4.4.1 NumMissing(d, l, k)

This heuristic states that the number of missing packets from device d in any sub mutation traces of length l shall not exceed k ($k \leq l$). The sliding window of size l serves two purposes. First, l should be large enough for the calculated packet loss ratio to be statistically meaningful. Second, it ensures that the packet losses are evenly distributed among the entire packet trace.

The intuition behind this heuristic is that the sniffer's empirical packet loss probability can usually be measured before validation. Therefore, the likelihood that the sniffer misses more packets than prior measured loss ratio is quite low. The value of l and k can then be configured such that k/l is marginally larger than the measured ratio.

4.4.2 GoBack(k)

This heuristic states that the search should only backtrack at most k steps when the search gets stuck using only E. The motivation is that many protocols operate as a sequence of independent transactions, and the uncertainty of previous transactions often does not affect the next transaction. For instance, in 802.11 packet transmission protocol, each packet exchange, include the original, retransmission and acknowledgment packets, constitutes a transaction. And the retransmission status of previous packets has no effect on the packets with subsequent sequence numbers, hence need not be explored when resolving the uncertainty of the packets with new sequence numbers. Note that we do not require the protocol to specify an exact transaction boundary, but only need k to be sufficiently large to cover a transaction.

5 Case Studies

We present case studies on applying our validation framework on two protocols implemented in the NS-3 network simulator: 802.11 data transmission and ARF rate control algorithm. The goal is to demonstrate how our framework can avoid false alarms and report true violations on incomplete sniffer traces and report true violations. We also report the application of our framework to a commercial product under development.

5.1 802.11 Data Transmission

In this section, we first show that our framework can improve validation precision by inferring the missing or extra packets using the augmented transition framework. We then demonstrate the ability of our framework to detect true violations by manually introducing bugs in the NS-3 implementation and show the precision and recall of validation results.

5.1.1 Experimental Setup

We set up two Wi-Fi devices acting as the transmitter (DUT) and receiver respectively. Another Wi-Fi device is configured in monitor mode and acts as the sniffer. During the experiments,we collect both the DUT packet trace (the ground truth) and the sniffer trace.

The DUT and endpoint are configured to use the IEEE 802.11g standard with both RTS/CTS and fragmentation disabled. A Constant Bit Rate (CBR) UDP traffic (54 Mbps) is generated from the DUT to the endpoint. The UDP packet size is 1436 bytes, which results in a 1500 bytes Wi-Fi packet.

In order to control the packet loss ratios between each pair of devices, we developed a new propagation loss model for NS-3 called `MatrixRandomPropagationLossModel`. Instead of a constant propagation loss as in existing `MatrixPropagationLossModel`, the signal propagation loss between a pair of nodes is determined by a binary random variable of two values: 0 dB (no loss) and 1000 dB (complete loss). Therefore, any packet loss probability can be achieved by adjusting the random variable distribution. The model supports both symmetric and asymmetric propagation losses. We use symmetric propagation loss in all our experiments. Finally, pcap capture is enabled in both the DUT and the sniffer devices.

5.1.2 Verifying Unmodified Implementation

In the original monitor state machine shown in Fig. 1, we set acknowledgment timeout $T_o = 334\mu$s, maximum retransmission delay $T_m = 15$ms according to

the protocol. We also adapt the state machine to include multiple retransmissions[2] instead of one.

Let Pr_{ds}, Pr_{es} and Pr_{ed} be the packet loss probability between the DUT and sniffer, endpoint and sniffer, endpoint and DUT respectively. Pr_{ed} represents the characteristics of the system being tested, while Pr_{ds} and Pr_{es} represent the sniffer's quality in capturing packets.

We vary each of the three probabilities, Pr_{ds}, Pr_{es} and Pr_{ed}, from 0 to 0.5 (both inclusive) with 0.05 step. For each loss ratio combination, we ran the experiment 5 times, and each run lasted 30 seconds. In total, 6655 ($11^3 \times 5$) pairs of DUT and sniffer packet traces were collected.

To establish the ground truth of violations, we first verify the DUT packet traces using the *original* state machine S. This can be achieved by disabling augmented transitions in our framework. As expected, no violation is detected in any DUT packet traces.

We then verify the sniffer traces using the augmented state machine S^+. For the $GoBack(k)$ heuristic, we set $k = 7$, which is the maximum number of transmissions of a single packet. For the $NumMissing(d, l, k)$ heuristic, we set the sliding window size $l = 100$, and $k = 80$ such that no violation is reported. The relationship of k and validation precision is studied in next section.

Next, we present detailed analysis of the augmented transitions on the sniffer traces. The goal is to study for a given system packet loss probability Pr_{ed}, how the sniffer packet loss properties (Pr_{ds} and Pr_{es}) affect the difference between the DUT trace and the mutation trace, which represents a guess of the DUT trace by the augmented state machine based on the sniffer trace.

For all following analysis, we divide the traces into three groups according to Pr_{ed}: low ($0 \leq Pr_{ed} \leq 0.15$), medium ($0.20 \leq Pr_{ed} \leq 0.35$) and high ($0.40 \leq Pr_{ed} \leq 0.50$).

The difference between two packet traces can be quantified by the Jaccard distance metric.

$$Jaccard(Tr_1, Tr_2) = \frac{|Tr_1 \ominus Tr_2|}{|Tr_1 \cup Tr_2|} \tag{1}$$

where $\ominus$ is the symmetric difference operator. The distance is 0 if the two traces are identical, and is 1 when the two traces are completely different. The smaller the distance is, the more similar the two traces are.

A naïve way to calculate the Jaccard distance of two traces is to use the hash of the $(time, packet)$ pair for set intersection and union operation. However, it does not work for mutation trace which contains fabricated packets with no actual payload. Therefore, we use a protocol specific canonical representation of packets when calculating the distance. In particular, the string $r\_DATA\_i\_t$ represents the t^{th} transmission of a data packet with sequence number i, and r represents the round of sequence numbers as it wraps after 4096. And similarly $r\_ACK\_i\_t$ is the corresponding acknowledgment packet.

[2] The exact number of retransmissions is not part of the protocol, and NS-3 implementation set this to be 7.

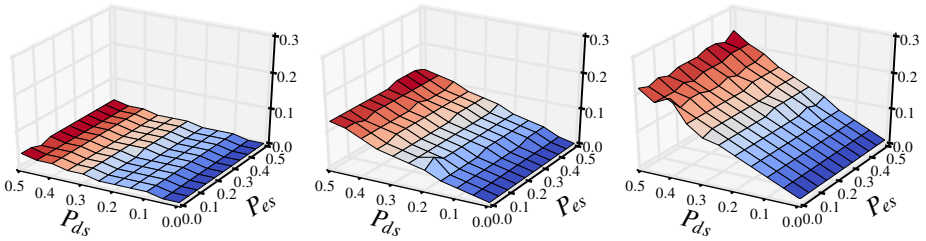

(a) $0.05 \leq Pr_{ed} \leq 0.15$ (b) $0.2 \leq Pr_{ed} \leq 0.35$ (c) $0.4 \leq Pr_{ed} \leq 0.5$

Fig. 5: **Jaccard Distance Between Mutation and DUT Traces.** For each data point, the mean of the 5 runs is used.

Fig. 5 shows the Jaccard Distance between mutation and its corresponding DUT trace. We make the following observations. First, for a given system loss probability Pr_{ed} (each sub-figure), the lower the sniffer packet loss probability (Pr_{ds} and Pr_{es}), the smaller Jaccard distance between the DUT and mutation trace. Intuitively, this means a sniffer that misses less packets can enable our framework to better reconstruct the DUT trace.

Second, we observe a *protocol-specific* trend that Pr_{ds} is more dominant than Pr_{es}. This is because retransmission packets of the same sequence number are identical, hence when the sniffer misses multiple retransmission packets, our framework only needs to infer one retransmission packet to continue state machine execution.

Finally, as the system loss probability Pr_{ed} increases, the Jaccard distance increases more rapidly as Pr_{ds} increases. This is because the ratio of retransmission packet increases along with Pr_{ed}.

We then evaluate the cost of resolving uncertainty. In particular, we use the Average Search Steps Per Packet (ASSPP) as a metric to quantify the search cost. It is calculated by dividing the total number of search steps by the number of packets in the packet trace. For DUT traces, ASSPP is always 1 since there is no uncertainty. For sniffer traces, however, multiple search steps must be conducted to resolve the potential uncertainty of each packet in the sniffer trace.

Fig. 6 shows the distribution of ASSPP at different Pr_{ed}. Similar to the case in Fig. 5, Pr_{ds} plays a dominant role in determine the searching cost. One interesting observation for this particular protocol is that the search cost peaks when Pr_{ds} is high while Pr_{es} is low. In such loss probability combinations, sniffer misses many data packets from the DUT but picks up lots of *dangling Ack* packets from the DUT. Because the *Ack* packet has neither sequence numbers nor retry flag, the searching algorithm had a hard time resolving such uncertainty.

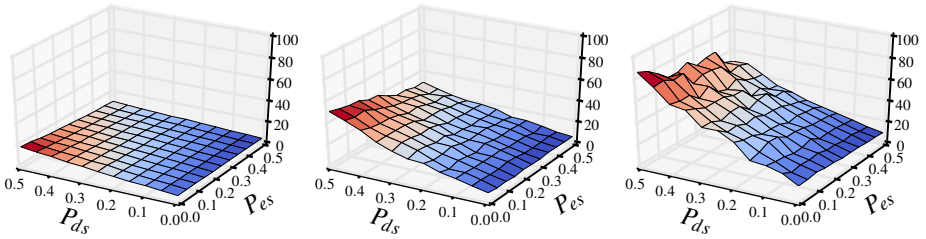

(a) $Pr_{ed} \in [0.05, 0.10, 0.15]$ (b) $Pr_{ed} \in [0.2, 0.25, 0.3, 0.35]$ (c) $Pr_{ed} \in [0.40, 0.45, 0.50]$

Fig. 6: **Average Searching Step Per Packet.** For each data point, the mean of 5 runs is used.

5.1.3 Introducing Bugs

We have demonstrated that our framework can tolerate sniffer imperfections and avoid raising false alarms. The next question is, can it detect true violations? To answer this question, we manually introduce several bugs in NS-3 implementation that concerns various aspects of 802.11 data transmission protocol. More specifically, the bugs are:

- **Sequence Number**: the DUT does not assign sequence number correctly. For example, it may increase sequence by 2 instead of 1, or it does not increase sequence number after certain packet, etc. We choose one type of such bugs in each run.
- **Semantic**: the DUT may retransmit even after receiving *Ack*, or does not retransmit when not receiving *Ack*.

We instrument the NS-3 implementation to embed instances of bugs in each category. At each experiment run, we randomly decide whether and which bug to introduce for each category. We fix $Pr_{ds} = Pr_{es} = 0.1$ and vary Pr_{ed} from 0.0 to 0.5 with 0.01 step. For each Pr_{ed} value, we ran the experiment 100 times, of which roughly 75 experiments contained bugs. In total, 5100 pairs of DUT and sniffer traces were collected.

We use the DUT packet traces as ground truth of whether or not each experiment run contains bugs. For each Pr_{ed} value, we calculate the precision and recall of violation detection using the sniffer traces.

$$\text{Precision} = \frac{|\{\text{Reported Bugs}\} \cap \{\text{True Bugs}\}|}{|\{\text{Reported Bugs}\}|} \tag{2}$$

$$\text{Recall} = \frac{|\{\text{Reported Bugs}\} \cap \{\text{True Bugs}\}|}{|\{\text{True Bugs}\}|} \tag{3}$$

The precision metric quantifies how *useful* the validation results are , while the recall metric measures how *complete* the validation results are.

Fig. 7 shows the CDF of precision and recall of the 51 experiments for various k values. For precision, as expected, the more tolerant the search to

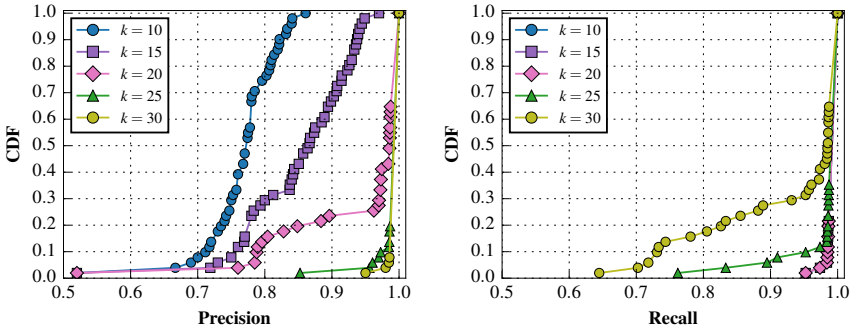

Fig. 7: **Precision and Recall of Validation Results.**

sniffer losses (larger k), the more tolerant the framework is to sniffer losses, and the more precise the violation detection. In particular, when $k = 30$, the precisions are 100% for all Pr_{ed} values. Second, the recall is less sensitive to the choice of k. Except for the extreme case when $k = 30$, all other thresholds can report almost all the violations.

5.2 ARF Rate Control Algorithm

We report a bug found in NS-3 ARF [21] implementation which causes the sender to get stuck at a lower rate even after enough number of consecutive successes. The bug was detected using sniffer traces and confirmed by both the DUT trace and source code inspection.

Automatic Rate Fallback (ARF) [21] is the first rate control algorithm in literature. In ARF, the sender increases the bit rate after Th_1 number of consecutive successes or Th_2 number of packets with at most one retransmission. The sender decreases bit rate after two consecutive packet failures or if the first packet sent after rate increase (commonly referred as *probing* packet) fails.

Fig. 8 shows the state machine S for the packet trace collected at sender (DUT), where $DATA_i^r$ denotes a data packet with sequence number i and bit rate r, $DATA_i^{r\prime}$ is a retransmission packet and Ack is the acknowledgment packet. The `pkg_succ` function is shown in Algorithm 3.

The `succ` variable is used to track the number of consecutive packet successes. It is increased after each packet success , and is reset to 0 after a rate increase or upon a packet failure ($s_1 \rightarrow s_2$). Similarly, `count` is to track the number of packets with at most one retransmission, and is increased after packet success, or for the first packet retransmission ($s_1 \rightarrow s_2$). It is reset when there are two consecutive packet failures ($s_2 \rightarrow s_3$). Finally, the `probe` flag is set upon rate increases to indicate the probing packet, and is cleared upon packet success. The variable `r` is the current bit rate, which is decreased if the probing packet fails ($s_1 \rightarrow s_4$), or every two consecutive failures ($s_2 \rightarrow s_3$). If

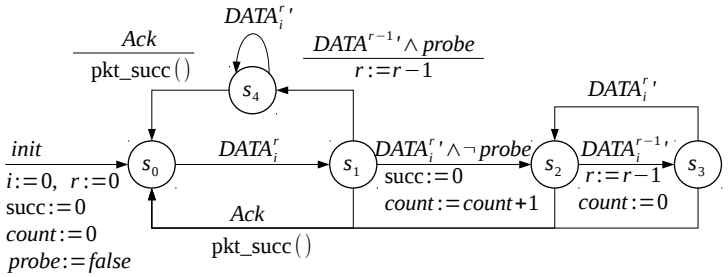

Fig. 8: **Monitor State Machine for ARF Rate Control Algorithm.** Timing constraints are omitted for succinctness.

Algorithm 3 `pkt_succ` function

```
1: function PKT_SUCC
2:     i := (i+1)%N
3:     succ := succ + 1
4:     count := count + 1
5:     probe := false
6:     if r < R and (succ ≥ Th1 or count ≥ Th2) then
7:         r := r+1
8:         succ := 0
9:         count := 0
10:        probe := true
```

`r` is not the highest rate, it is increased when either of the two thresholds are reached.

In particular, the bug we found lies in the implementation's `pkt_succ` function in line 6. Instead of checking `count` $\geq$ `Th_2`, the implementation checks `count == Th_2`. This bug also exists in the NS-3 implementation of Adaptive ARF (AARF) algorithm [22] and the pseudo code implementation of AARF in [23].

Note that the `count` variable is incremented twice if a packet succeed after one retransmission: once in $s_1 \rightarrow s_2$, once in the `pkt_succ` function for the retransmission packet. Therefore, if the value of `count` is $Th_2 - 1$ and the next packet succeed after one retransmission, the value of `count` will be $Th_2 + 1$, which would fail the implementation's test of `count == Th_2`.

5.3 Industrial Application

We have applied our framework for runtime verification of the wireless protocol implementation of a commercial product that has several millions of shipping devices. The product runs a proprietary wireless protocol to reduce latency and power consumption. Sniffer traces were collected in regular testing process, but were only manually inspected previously.

We obtained 75 sniffer traces from the testing team for a new version of the protocol that is under development and testing. This team has been testing the

Protocol Aspects	Traces	Violations (%)
Sequence Number	3049	1539 (50.48%)
Station Scheduling	3046	2045 (67.14%)
Uplink Modulation	3127	8 (0.26%)
Downlink Modulation	3127	24 (0.77%)

Table 1: **Validation Results on Traces from the Gaming Controller Wireless Protocol.**

implementation for a few weeks. Each trace contained around 6 million packets that were captured during 1 hour and 40 minutes of protocol operation.

We first split the traces into 100,000 packet segments, which yields 3127 traces for testing. We then applied our framework on the traces to validate the DUT implementation. We found that the latest implementation of the protocol under development had violations of the protocol specification. Some of the implementation bugs we found related to the sequence number management, station scheduling during power saving mode, and modulation rates adaptation. Table 1 summarizes the validation results. The **Traces** column shows the number of traces that we have successfully validated for each category, and the third column shows the fraction of traces that contain at least one violation in that category.

Note that if we disable the augmented transitions, each trace will be flagged as violation because of the missing packets, thereby reducing the usability of the tool. We also note that some bugs manifest more often than others. For instance, the bugs related to packet sequence number and station scheduler were detected in about half the traces, and the bug related to the rate control algorithm was detected in only a few traces. This is because the previous two aspects are essential in all protocol operations, while the bugs related to rate control only manifest themselves under certain link conditions.

After communicating with the testing team, we confirmed that the sequence number bug was already known, as it is relatively easy to detect even by manually examining the traces. The bug related to station scheduling was also noticed before, yet no quantitative results about how frequent this bug appears were obtained because of the lack of automatized validating tools. Finally, the bug related to rate control, which was unknown previously, has been filed as a bug report. All reported bugs were fixed before the next release the product.

6 Related Work

Hidden Markov Model (HMM) Approach. When considering the whole system under test (both DUT and endpoint), the sniffer only captures a subset of the all the packets (events). This is similar to the event sampling problem in runtime verification [6,18,2,14,5,33]. In particular, Stoller *et al* [34] used HMM-based state estimation techniques to calculate the confidence that the

temporal property is satisfied in the presence of gaps in observation. HMM-based approaches proposed in [34] suffer from overhead of both running time and memory consumption. Optimizations such as approximate precomputed tables [4] and particle filtering [20] were later proposed.

In our problem, there are not only gaps in the observation (packets missed by the sniffer), but also uncertainty regarding observed events, which lies in whether a packet in the sniffer trace was received by its destination device. This unique challenge makes it infeasible to directly apply the HMM-based approaches proposed in [34]. Besides, there are several advantages of our proposed augmented transition approach. First, the automatically augmented state machine precisely encodes the protocol specification and the uncertainty. This is intuitive to design and natural for reporting the evidence for a trace being successful. We do not require a user to specify the number of states of the underlying HMM, or accurately provide underlying probabilities. Second, we use timed automata to monitor the timing constraints which are common in wireless protocols. It may be non-trivial to encode such timing information in HMM. Finally, we can exploit domain knowledge to devise effective pruning heuristics to rule out unlikely sequences during the exhaustive search.

Edit Distance. The closet work to ours is [19], which uses the weighted edit distance to measure the robustness between digitized Cyber-Physical System (CPS) signals and STL specifications. The automaton of verifying STL properties is first translated into weighted edit automaton, which explicitly handles substitution, insertion and deletion by augmenting the original automaton with transitions and associating to them the appropriate weight function. An observed signal sequence is then input into the weighted edit automaton to compute the weighted edit distance between the observation and the specification.

We share the same spirit with [19] in that we also explicitly handle the sniffer trace uncertainty using the augmented transitions. However, there is no need for substitution in our application scenario, as the sniffer can only either receive or miss a packet, but can not mis-interpret packets. In addition, [19] only considers real value observations, and relies on this restriction to define the edit distance for insertion transitions. The approach can not be directly applied in our case since we need to deal with compound observations (packets).

Network Protocol Validation. Lee *et al* [24] studied the problem of passive network testing of network management. The system input/output behavior is only partially observable. However, the uncertainty only lies in missing events in the observation, while in the context of wireless protocol verification, the uncertainty could also be caused by extra events not observed by the tested system. Additionally, they do not provide any formal guarantees even for cases when we report a definite bug. Software model checking techniques [27,16] have also been used to verify network protocols. Our problem is unique because of the observation uncertainty caused by sniffers. Our framework shares similarity with *angelic verification* [10] where the program verifier reports a warning only when no acceptable specification exists on unknowns.

Testing Under Uncertainty. The position paper by Rosemblum *et al* [13] contains excellent motivation for the need to combat uncertainty foundationally when testing systems. McKinley *et al* [7,30] deals with checking assertions in programs dealing with noisy data from sensors. Instead of checking the truth or falsity of assertions, they model the probability distribution of the assertion conditions and perform Monte-carlo based simulations to estimate the probabilities. Our work can be seen as leveraging non-determinism to weaken the specification logically to precisely define the problem complexity, and use probabilities to guide the search for likely mutations. Other works have used sampling to find data-race bugs [26], and ensure that the sampling does not lead to spurious alarms.

7 Conclusions

We formally define an instance of the uncertainty problem in validating wireless protocol implementations using sniffers. We describe a systematic augmentation of the protocol state machine to explicitly encode the uncertainty of sniffer traces. We characterize the NP-completeness of the problem and propose both an exhaustive search algorithm and heuristics to restrict the search to more likely traces. We present two case studies using NS-3 network simulator to demonstrate how our framework can improve validation precision and detect real bugs. We also report the application of our framework on a commercial product under development.

Finally, we discuss a few challenges and future directions.

Verification Coverage. Given a single sniffer trace, it is possible that not all the states in the state machine are visited during the verification process. For instance, a rate control state machine based on certain consecutive packet losses patterns can not be verified if no such consecutive losses appear in the sniffer trace. In general, given a protocol state machine, it is challenging to extract the packet patterns for each state to be reached and to alter the testing such that such patterns can be observed.

Automated State Machine Construction. We manually constructed the protocol monitor state machines for the protocols studied in this paper based on the source code, comments and documentation. The process involves lots of labor effort and is time-consuming. A potential alternative is to automatically learn the sketch of the monitor state machines from the sniffer traces. Domain knowledge can then be leveraged to improve the sketch state machines.

References

1. R. Alur and D. L. Dill. A theory of timed automata. *Theoretical computer science*, 126(2):183–235, 1994.
2. M. Arnold, M. Vechev, and E. Yahav. Qvm: an efficient runtime for detecting defects in deployed systems. In *ACM Sigplan Notices*, volume 43, pages 143–162. ACM, 2008.

3. P. Bahl, R. Chandra, J. Padhye, L. Ravindranath, M. Singh, A. Wolman, and B. Zill. Enhancing the security of corporate Wi-Fi networks using DAIR. In *Proceedings of the 4th international conference on Mobile systems, applications and services*, pages 1–14. ACM, 2006.
4. E. Bartocci, R. Grosu, A. Karmarkar, S. A. Smolka, S. D. Stoller, E. Zadok, and J. Seyster. Adaptive runtime verification. In *International Conference on Runtime Verification*, pages 168–182. Springer, 2012.
5. D. Basin, F. Klaedtke, S. Marinovic, and E. Zălinescu. Monitoring compliance policies over incomplete and disagreeing logs. In *International Conference on Runtime Verification*, pages 151–167. Springer, 2012.
6. B. Bonakdarpour, S. Navabpour, and S. Fischmeister. Sampling-based runtime verification. In *FM 2011: Formal Methods*, pages 88–102. Springer, 2011.
7. J. Bornholt, T. Mytkowicz, and K. S. McKinley. Uncertain¡ t¿: A first-order type for uncertain data. *ACM SIGARCH Computer Architecture News*, 42(1):51–66, 2014.
8. Y.-C. Cheng, J. Bellardo, P. Benkö, A. C. Snoeren, G. M. Voelker, and S. Savage. *Jigsaw: solving the puzzle of enterprise 802.11 analysis*, volume 36. ACM, 2006.
9. M. Ciabarra. WiFried: iOS 8 WiFi Issue. `https://goo.gl/KtRDqk`.
10. A. Das, S. K. Lahiri, A. Lal, and Y. Li. Angelic verification: Precise verification modulo unknowns. In *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I*, pages 324–342, 2015.
11. digitalmediaphile. Windows 10 wifi issues with surface pro 3 and surface 3. `http://goo.gl/vBqiEo`.
12. S. Edelkamp, V. Schuppan, D. Bošnački, A. Wijs, A. Fehnker, and H. Aljazzar. Survey on directed model checking. In *International Workshop on Model Checking and Artificial Intelligence*, pages 65–89. Springer, 2008.
13. S. Elbaum and D. S. Rosenblum. Known unknowns: Testing in the presence of uncertainty. In *Proceedings of the 22Nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, FSE 2014, pages 833–836, New York, NY, USA, 2014. ACM.
14. L. Fei and S. P. Midkiff. Artemis: Practical runtime monitoring of applications for execution anomalies. In *ACM SIGPLAN Notices*, volume 41, pages 84–95. ACM, 2006.
15. Gizmodo. The worst bugs in android 5.0 lollipop and how to fix them. `http://goo.gl/akDcvA`.
16. P. Godefroid. Model checking for programming languages using verisoft. In *Proceedings of the 24th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 174–186. ACM, 1997.
17. Google. Google contact lens. `https://en.wikipedia.org/wiki/Google_Contact_Lens`.
18. M. Hauswirth and T. M. Chilimbi. Low-overhead memory leak detection using adaptive statistical profiling. In *Acm Sigplan Notices*, volume 39, pages 156–164. ACM, 2004.
19. S. Jakšić, E. Bartocci, R. Grosu, and D. Ničković. Quantitative monitoring of stl with edit distance. In *International Conference on Runtime Verification*, pages 201–218. Springer, 2016.
20. K. Kalajdzic, E. Bartocci, S. A. Smolka, S. D. Stoller, and R. Grosu. Runtime verification with particle filtering. In *International Conference on Runtime Verification*, pages 149–166. Springer, 2013.
21. A. Kamerman and L. Monteban. Wavelan®-ii: a high-performance wireless lan for the unlicensed band. *Bell Labs technical journal*, 2(3):118–133, 1997.
22. M. Lacage, M. H. Manshaei, and T. Turletti. IEEE 802.11 rate adaptation: a practical approach. In *Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*, pages 126–134. ACM, 2004.
23. M. Lacage, M. H. Manshaei, and T. Turletti. IEEE 802.11 rate adaptation: a practical approach. *[Research Report] RR-5208*, (¡inria-00070784¿):25, 2004.
24. D. Lee, A. N. Netravali, K. K. Sabnani, B. Sugla, and A. John. Passive testing and applications to network management. In *Network Protocols, 1997. Proceedings., 1997 International Conference on*, pages 113–122. IEEE, 1997.
25. R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Analyzing the mac-level behavior of wireless networks in the wild. In *ACM SIGCOMM Computer Communication Review*, volume 36, pages 75–86. ACM, 2006.

26. D. Marino, M. Musuvathi, and S. Narayanasamy. Literace: effective sampling for lightweight data-race detection. In *ACM Sigplan Notices*, volume 44, pages 134–143. ACM, 2009.
27. M. Musuvathi, D. Y. Park, A. Chou, D. R. Engler, and D. L. Dill. CMC: A pragmatic approach to model checking real code. *ACM SIGOPS Operating Systems Review*, 36(SI):75–88, 2002.
28. T. Mytkowicz, P. F. Sweeney, M. Hauswirth, and A. Diwan. Observer effect and measurement bias in performance analysis. 2008.
29. G. F. Riley and T. R. Henderson. The ns-3 network simulator. In *Modeling and Tools for Network Simulation*, pages 15–34. Springer, 2010.
30. A. Sampson, P. Panchekha, T. Mytkowicz, K. S. McKinley, D. Grossman, and L. Ceze. Expressing and verifying probabilistic assertions. In *ACM SIGPLAN Notices*, volume 49, pages 112–122. ACM, 2014.
31. Savvius Inc. Savvius Wi-Fi adapters. `https://goo.gl/l3VXSx`.
32. J. Shi, S. K. Lahiri, R. Chandra, and G. Challen. *Wireless Protocol Validation Under Uncertainty*, pages 351–367. Springer International Publishing, Cham, 2016.
33. A. P. Sistla, M. Žefran, and Y. Feng. Runtime monitoring of stochastic cyber-physical systems with hybrid state. In *International Conference on Runtime Verification*, pages 276–293. Springer, 2011.
34. S. D. Stoller, E. Bartocci, J. Seyster, R. Grosu, K. Havelund, S. A. Smolka, and E. Zadok. Runtime verification with state estimation. In *Runtime Verification*, pages 193–207. Springer, 2011.
35. Wikipedia. Chromecast. `https://en.wikipedia.org/wiki/Chromecast`.
36. Wikipedia. Xbox One controller. `https://en.wikipedia.org/wiki/Xbox_One_Controller`.