







Network Simulator-centric Compositional Testing

Tom Rousseaux , Christophe Crochet , John Aoga , and Axel Legay 

INGI, ICTEAM, Université catholique de Louvain, Place Sainte Barbe 2, L05.02.01,
1348 Louvain-La-Neuve, Belgium

{firstname.lastname}@uclouvain.be

Abstract. This article introduces a novel methodology, Network Simulator-centric Compositional Testing (NSCT), to enhance the verification of network protocols with a particular focus on time-varying network properties. NSCT follows a Model-Based Testing (MBT) approach. These approaches usually struggle to test and represent time-varying network properties. NSCT also aims to achieve more accurate and reproducible protocol testing. It is implemented using the Ivy tool and the Shadow network simulator. This enables online debugging of real protocol implementations. A case study on an implementation of QUIC (*picoquic*) is presented, revealing an error in its compliance with a time-varying specification. This error has subsequently been rectified, highlighting NSCT’s effectiveness in uncovering and addressing real-world protocol implementation issues. The article underscores NSCT’s potential in advancing protocol testing methodologies, offering a notable contribution to the field of network protocol verification.

Keywords: Model-Based Testing, Time-varying Network Properties, Software verification and validation, Formal Specifications, Network Simulator, Internet protocols, QUIC, Concrete Implementation, Adverse Stimuli

1 Introduction

Ensuring the safety and effectiveness of systems is paramount. One way to achieve this goal is through the use of model-checking approaches. These approaches employ mathematical models of the system and exhaustively check the specifications against all possible behaviors of the system. Examples of such approaches include *SPIN* [27, 42] and *NUSMV* [20, 43], which use Linear-Temporal Logic (LTL) [53] or Computation Tree Logic (CTL) [18] to describe specifications. First, model checking results were applied to mathematical models of the system under validation. However, over the last decades, we have seen the emergence of techniques applied directly to implementations [19]. An inherent hurdle in model-checking lies in the state-space explosion dilemma triggered by exhaustive exploration of the entire state-space.

To tackle this challenge, researchers have proposed Statistical Model Checking (LLTYSG19, LL20). This approach entails simulating the system and using statistical algorithms to ascertain whether it meets a measurable speci-

fication within a finite execution with a certain probability and given confidence. The approach, which has been implemented in tools such as *UPPAAL-SMC* [22,34,35,37] and *PLASMA* [9,13,39], has been applied on a wide range of case studies [36]. Statistical Model Checking is primarily employed for validating properties on mathematical models. Except for specific instances like in [49], direct validation on code has been rarely proposed. Generating fair traces from the code required by the statistical algorithm is challenging. Furthermore, with few exceptions as seen in [17], specifications are typically formulated using Bounded LTL. This representation is inadequate for describing the specifications of a network protocol. In this paper, we propose an approach that leverages the principle of simulation akin to Statistical Model Checking, but within the framework of specifications described in a sophisticated language and directly validated against the implementation. This approach, known as model-based testing [50], offers a more scalable solution. Our approach also introduces specific techniques to address the challenges involved in protocol verification.

In protocol verification, traditional methodologies rely on multiple independent implementations and interoperability testing to validate protocol designs. However, comprehensive model-based verification is often lacking in well-known approaches. Ivy, a notable exception, allows working with protocol implementations and adversarial stimuli.

Ivy’s mathematical model [51] serves as the language describing the system specifications, whereas model-checking approaches typically involve two mathematical models, one for the system and one for the specifications.

Network-centric Compositional Testing (NCT) [45] is an emerging methodology that was introduced within Ivy. NCT introduces a formal statement of a protocol standard, allowing effective testing of implementations for compliance, not just interoperability. NCT uses formal specifications of protocols to automatically create testing tools. These tools generate random test cases by solving constraints with the help of an SMT solver. This enables adversarial testing in real-world environments, uncovering compliance issues and ambiguities in standard protocol documents (RFCs). NCT uncovered errors and vulnerabilities in the real-world protocol QUIC [31], proving its effectiveness.

Although NCT serves as a foundation for network-centric protocol verification, it does not address time-varying network properties. Time-varying network properties describe the timed aspects of network protocols. These properties include internal timeouts to, for example, trigger a packet retransmission. These are involved in properties that are more complex to model, such as congestion control schemes in retransmission mechanisms.

Ivy deterministically generates output packets from input packets, but does not provide computation time requirements. This duration can non-deterministically exceed (or not) protocol timeouts. This can impact the inputs, for example by triggering a packet retransmission. This non-determinism prevents the reproducibility of the experiments. If Ivy discovers a bug in an implementation, Ivy is not necessarily able to reproduce it.

We have developed a new approach called *Network Simulator-centric Compositional Testing* (NSCT) to address these limitations. NSCT is designed to focus on verifying time-varying network properties in network protocols and ensuring experimental reproducibility.

Our method extends Ivy to support time-related features, providing a more network-centric approach. Additionally, we have integrated the Shadow network simulator which allows online debugging of protocol implementations. This integration ensures the determinism and reproducibility of the experiments. We have successfully applied our approach to test the QUIC protocol and have demonstrated its effectiveness in verifying time-varying network properties. NSCT reveals an error on the *picoquic* implementation and the QUIC idle timeout connection termination. This approach brings advancements to the network community by enabling more detailed and accurate protocol testing and verification. It enables specifying loss detection and congestion control in QUIC defined by RFC9002 [29].

The remainder of this paper is structured as follows. Section 2 provides background information on different verification types and the emergence of Ivy. Section 3 outlines our methodology, providing details about *Network Simulator-centric Compositional Testing* (NSCT). Section 4 delves into the verification of time-varying network properties in QUIC. Then, Section 5 discusses our findings and proposes avenues for future work. Finally, Section 6 leads to the related work and the conclusion.

2 Background

There are two main ways to create adversarial tests for network protocols: with [10, 52, 61] or without [3, 12, 16, 38, 54] checking compliance with a standard. Approaches that do not check compliance with a standard include fuzz testing [38], white-box testing [12, 54], and other methods that create a verified reference implementation [3] or prove properties of an existing implementation [16]. On the contrary, approaches that verify compliance with a standard, also known as Model-Based Testing (MBT) [50], involve constructing an abstract model with Finite State Machines (FSMs) to explore and generate test scenarios [10, 52, 61]. However, the incorporation of data into FSMs adds significant complexity and challenges to these formalisms [45].

Network-centric compositional (NCT) approaches avoid the use of FSMs. To grasp this and our approach, which builds upon NCT, it is crucial to understand the functioning of the *Ivy tool* that implements them.

Ivy is a verification tool implementing multiple proving techniques [44, 51]. It is used to correct the design and implementation of algorithms and distributed protocols. It supports modular specifications and implementation. Ivy is used to interactively verify the safety properties of infinite-state systems. Ivy introduced a Relational Modeling Language (RML). This language allows describing the state of a program using formulas from first-order logic and uses relations

(boolean predicate), functions, modules, and type objects as the main abstractions to represent the state of the system. Let us illustrate the functioning *Ivy* using a running protocol example.

MiniP (Minimalist Protocol) is a simple protocol. MiniP defines packets that contain frames. Any packet must contain exactly two frames. Three types of frames are defined: PING, PONG, and TIMESTAMP. PING frame contains a four-byte string representing the word "ping". PONG frame also contains a four-byte string expressing the word "pong". The PING frame or the PONG frame must be present in a packet. Finally, the TIMESTAMP frame contains an eight-byte unsigned integer representing the moment, in milliseconds, when the packet is sent. This frame must be present in all packets.

Figure 1 represents the finite state machines (FSM) of MiniP. The client starts by sending a packet containing the PING frame followed by the TIMESTAMP frame as payload. The server must then respond within three seconds, with a packet containing the PONG frame followed by the TIMESTAMP frame. This exchange continues until the client stops the connection. The client terminates the connection by not transmitting any packets for more than three seconds.

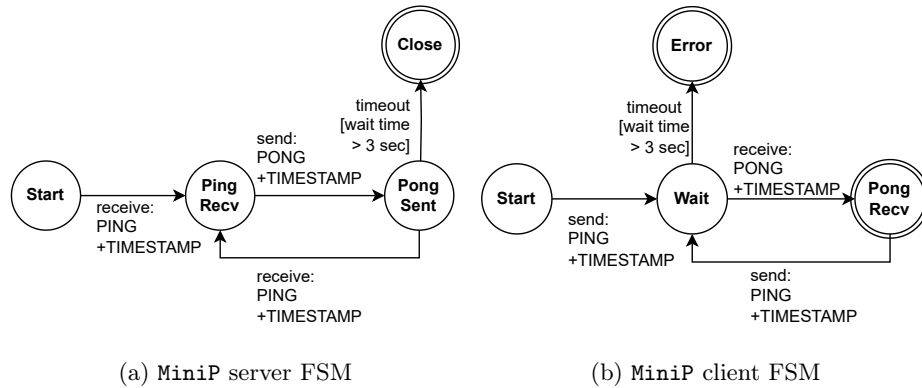


Fig. 1: MiniP Finite state machines (FSM)

Some MiniP components implementation with Ivy The first and most important components to implement for this protocol are frames. In Ivy, a frame is implemented using the *type object*. Listing 1 provides an example of the frame object, which includes a subtype object representing the PING frame and a generic action `handle(f:frame)` that must be implemented by subtype objects. The PING frame defines a `data` field containing the four-byte payload as described in the specification.

1. State example

```

1  object frame = {
2    type this
3    object ping = {
4      variant this of frame = struct {
5        data : stream_data
6      }
7    }
8    action handle(f:this) = {
9      require false;
10   }
11 }

```

Ivy's "action" statement is used to manipulate the states and add requirements. An action can be considered as a procedure and cannot be stored in variables or passed as arguments. Listing 2 illustrates the `handle(f:frame)` action, which encompasses all the properties associated with the PING frame and adds requirements that will be checked every time a PING stream is received and generated. Lines 7 and 8 specify that the data payload must be a "ping" and have a length of four bytes. Line 9 requires that a PING frame should not be present in a packet using the relation `ping_frame_pending` defined on Line 1. Line 12 illustrates the invocation of the `enqueue_frame(f:frame)` action, which also modifies various states within the model and is used to append a frame to a packet object.

2. Object procedure example

```

1  relation ping_frame_pending
2
3  object frame = {
4    object ping = {
5      action handle(f:frame.ping)
6      around handle {
7        require f.data = ping_data;
8        require f.data.end = 4;
9        require ~ping_frame_pending;
10       ...
11       ping_frame_pending := true;
12       call enqueue_frame(f);
13     }
14   }
15 }

```

Network-centric Compositional Testing methodology (NCT) NCT, a specialized approach within Model-Based Testing (MBT), is specifically designed for network protocols. It provides a structured method for creating formal specifications of Internet protocols and subsequently testing them [45]. The NCT principle is demonstrated in Figure 2 using Ivy. The process begins with converting the RFC into an Ivy formal model (a). Once the Ivy code is parsed, a generator is used to create concrete and randomized testers (b). Finally, the implementation of the real-life protocol is tested and verified against the testers that employ an SMT solver to satisfy the constraints of the formal protocol requirements. When a requirement fails, the resulting traces (c) can be analyzed to identify any potential errors or vulnerabilities.

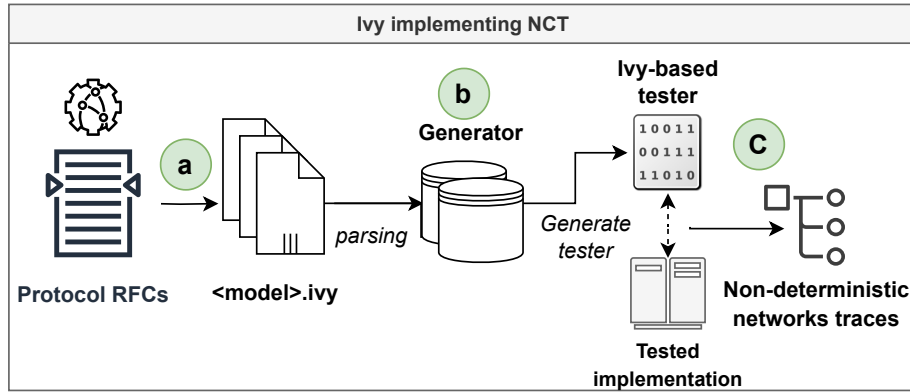


Fig. 2: Ivy implementing NCT

The NCT principle is based on compositional testing, which views formal specifications as a set of interconnected components/processes with their corresponding inputs and outputs. This approach allows testing protocol behaviors as observed on the wire, rather than relying on an abstract mathematical model of the protocol. This is why this methodology is called "network-centric".

The design of MiniP is in line with the principles of the NCT methodology. In MiniP, the "Frame" process produces output that serves as input for the "Packet" process. The "assumptions" regarding the inputs of a process are treated as "guarantees" for the outputs of other processes. Figure 3 provides a visual representation of this structure. In the context of MiniP, each element represents a layer of the MiniP stack, including the frame layer (a) and the packet layer (b). The *shim* component (c) is responsible for transmitting and receiving packets across the network. When a packet is received, the *shim* invokes the `ping_packet_event` action. This action contains all the specifications associated with the MiniP packet and will generate an error if any of the requirements are not met. For instance, it verifies that a packet always contains two frames in the correct order. The frames are similarly managed with their respective actions. In Figure 3, the set of requirements is connected to the packet component (b).

Limitations NCT's success derives from how effectively it identifies errors and vulnerabilities in real-world protocols like QUIC [21, 46]. However, NCT presents some limitations associated with its inability to test time-varying network properties.

For instance, it cannot model the congestion control mechanism specified in RFC9002. Calculating the time needed for packet generation and verification of received packets can be significant, particularly when implementations under test generate packets in bursts, leading to false congestion due to increased round-trip time.

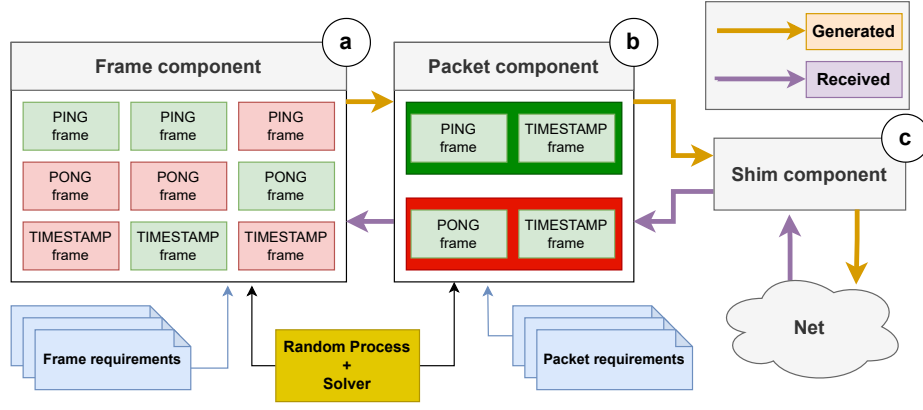


Fig. 3: MiniP Network-Centric Testing (NCT) structure

Additionally, debugging implementations is challenging because traces have no guarantee of reproducibility. In summary, here are the current main limitations of NCT for protocol testing:

- ① *Lack of the expressiveness to handle time-varying network properties [46].* NCT lacks the necessary capabilities to reason about precise time intervals and deadlines. It also does not provide guarantees on thread-scheduling or computation time. For example, NCT cannot verify whether a **MiniP** server will respond to a **PING** within three seconds.
- ② *Non-reproducible experiments.* While NCT offers deterministic verification for formal properties, this determinism does not extend to the network (*network nondeterminism*) or the implementations being tested (*internal nondeterminism*). As a result, the experiments cannot be reproduced. For example, a **MiniP** server that crashes when sending an odd **TIMESTAMP** will not crash deterministically when tested with NCT.
- ③ *Computational time exceeding protocol timeout.* The computation time required to verify incoming packets and generate packets that satisfy the model can interfere with the standard behavior defined by some protocols. This also hinders the reproducibility of the experiment. For instance, a **MiniP** server implementation in NCT may take too long to check the integrity of a **PING** and exceed the response time limit of 3 seconds. This issue would not occur with **MiniP** implementation on a modern computer, but it arises with real protocols due to the inherent complexity of their RFCs, as with **QUIC** [46].

3 Network Simulator-centric Compositional Testing

Network Simulator-Centric Compositional Testing (NSCT) is a specialized approach within Model-Based Testing (MBT) that aims to overcome the limitations of NCT discussed previously. NSCT, similar to NCT, adopts a network-centric perspective and employs a combination of two key ingredients.

Ingredient 1: Introduction of Network Simulators (NS)

Type of Network Simulators NS tools permit running a model or a real executable inside a controlled network environment. Model-oriented simulators are mainly used to verify protocols during their development stage [11]. Many types of NS exist; we will focus on time-dependent NS tools that have two main properties: they proceed chronologically and maintain a simulation clock [28]. This clock is essential to verify time-related properties. There are two types of *time-dependent* NS: time-driven and event-driven.

(1) *Time-driven NS* advance their clocks exactly by a fixed interval of δ time units [28]. This means that the simulation has a time precision of δ . To increase precision, δ must be small, which slows down the simulation computation.

(2) *Event-driven NS*, by comparison, advance their clocks by variable steps. Such tools progress as events unfold; each time step jumps precisely to the next event:

Event-driven network simulator [25]

```

while simulation is in progress do
  | remove smallest time-stamped event;
  | set time to this timestamp;
  | execute event handler;
end

```

Our approach using NS Our solution involves integrating NCT with an event-driven, time-dependent network simulator to effectively overcome the limitations ② (*Non-reproducible experiments*) and ③ (*Computational time*), and partially limitation ① (*Time-varying network properties*).

Addressing Limitation ②: Model-based testers and protocol implementations are real software components that must be executed rather than just modeled. When they are executed in a controlled environment, it becomes possible to stabilize and replicate desired random behaviors like encryption. This solution addresses the second limitation ②. This ensures the determinism of model-based protocol testing.

Addressing Limitation ③: Formal specifications developed to test internet protocols with respect to NCT focus on packet events. This means that the latency of the (network) link between the model-based test and the implementation under test (IUT) determines the clock steps. Assuming a latency of l , if the IUT sends a packet at time t , the model-based tester will receive it at time $t+l$. If the model-based tester responds immediately (without waiting for a specific delay to verify a timing property), the IUT will receive the response at time $t+2l$. This resolves ③ as the computation time does not affect the time perceived by the IUT.

Addressing (partially) Limitation ①: To address ① (*time-varying network properties*), it is essential to employ a time-dependent NS. However, simply using a time-dependent NS is not enough. It is also necessary to have a formal verification tool that can interact with the simulation clock. We will discuss this further in the second ingredient of NSCT.

Additional values using NS: The use of NS provides the ability to manage and control the network. It simplifies the creation of various network-related situations, including connection migration as described in RFC9000. An NS also enables realistic simulation scenarios of advanced modern network protocols [26]. In the following sections, we will introduce the specific NS that we used for NSCT.

Specific NS There are two modern discrete-event network simulators:

(1) *The ns-3 [56]* simulator is a freely available tool that has been specifically designed for research and educational purposes. It operates using models. To enable the execution of direct code within *ns-3*, the *DCE* framework [57] intercepts system calls and links them to *ns-3*. However, *DCE* does not support many of the necessary system calls required for protocol implementation, and the environment it supports has become outdated. The use of *ns-3 DCE* to simulate QUIC implementations requires significant effort and inhibits tool longevity. Researchers have recently expressed concerns about various challenges encountered while trying to simulate QUIC implementations using the *ns-3 DCE* framework [1].

(2) *Shadow [32, 33]* is a free, open source simulator. It was primarily designed to simulate Tor networks. Shadow works by intercepting a subset of the system calls (syscalls), simulating network calls. Despite lacking support for some key system calls initially, the Shadow project remained highly active and has since added built-in implementation for several important syscalls that were originally missing. This is a positive sign for the long-term sustainability of the tool.

Shadow provides a range of network-specific functionalities that are highly beneficial for researchers and engineers. It allows users to carefully design the network topology that they want to simulate, specifying nodes, links, and their connections. Shadow allows users to adjust parameters like link latency and

jitter, which are crucial for evaluating the performance of networked applications under different network conditions. Beyond its flexible design capabilities (e.g., configurable topologies), Shadow offers live debugging features to monitor and troubleshoot network behaviors in real-time during simulations.

Shadow easily supports single-threaded and multithreaded implementations. For example, if there is a multithreaded server connected to two clients, Shadow enables deterministic debugging of one client while allowing the other client and the server to operate independently within the simulation. This level of control and precision in debugging, even in complex multithreaded scenarios, provides researchers and developers with valuable insights into the behavior and interactions of various components. As a result, it improves the thoroughness of protocol analysis.

Ingredient 2: Integration of time-varying network properties testing in Ivy

In practice, network link properties are designed in the Ivy model or directly with Shadow. Ivy then builds the simulation configuration file with those properties and references the executable used for the test. Finally, Shadow launches the IUT and the Ivy test.

Adapting Ivy for Event-Driven Network Simulation Our approach aims to enhance the compatibility of the Ivy verifier with event-driven network simulators. To improve protocol verification in Ivy, we propose an adaptation that introduces an interface for manipulating time-related actions/relations. The interface is implemented in C++ and leverages the *'time.h'* library to facilitate the interception of system calls (syscalls) by the Shadow simulator.

This interface provides several key functionalities, including the manipulation of time in various units (seconds, milliseconds, microseconds), timer control (start and stop actions), and current-time querying. The interface supports setting time breakpoints at specific events and implements both blocking and non-blocking sleep mechanisms. Using non-blocking sleep allows the simulation to receive network events while "sleeping," resulting in simulations that more accurately mirror real-world network behavior.

Ivy's time interface is extensible, which allows for further adaptations and enhancements to meet the evolving needs of network protocol verification. The time interface is represented at Figure 4 (a).

Our progress in Ivy for simulating event-driven networks is built on an improved method for controlling event generation. We use signal handlers along with time-based signals such as **SIGALRM** to accurately manage event timing. This approach is illustrated in Figure 4 (b). It is especially effective in situations that require delayed responses, as it allows precise control over the timing of event generation and processing.

Finally, while Shadow's ability to modify network conditions is beneficial, it lacks flexibility, as it cannot vary the delay during the connection. To address this limitation, we developed a formal model that represents network quality,

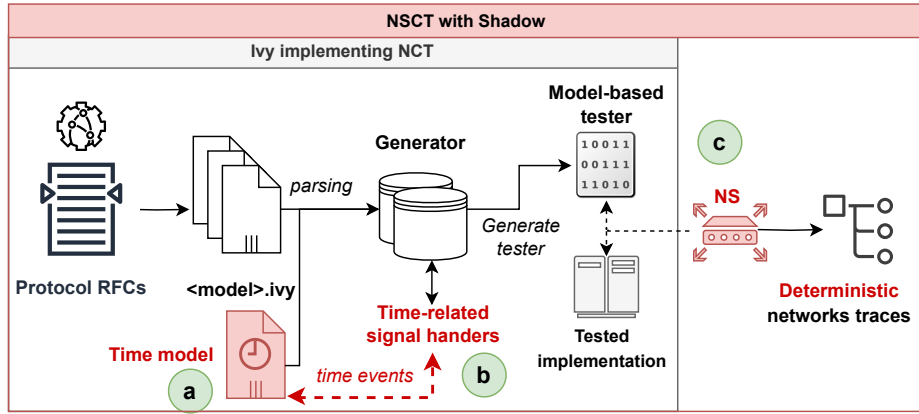


Fig. 4: Protocol Formal Verification toolchain (PFV)

enabling us to simulate more specific scenarios of network condition. Nevertheless, we still rely on Shadow for reproducibility and intercepting time syscalls as shown in Figure 4 ©.

Monitoring time-varying properties We can now use Shadow intercepting time syscalls to define safety properties for time-varying properties in Ivy without having to modify the tools directly. Using the implemented time module and the standard Ivy key words such as "require" or "assume", we can model all the first-order logic formula with time as variable or predicate.

Protocol Formal Verification (PFV) PFV¹ toolchain implements NSCT by leveraging Ivy and Shadow. The usability of the previous work is enhanced by implementing a multistage docker containerisation procedure [47], coupled with microservices and a basic graphical interface to initiate experiments. This architecture allows for easy testing of new protocol implementations with Ivy and Shadow. All containers implement a REST API to start Ivy.

Case Study - MiniP Protocol In our MiniP formal specification example, we implemented the property that the PONG message should be received within 3 seconds after the PING message being sent, as seen in Listing 3. To achieve that, we use the concept of time breakpoint. Then we add requirement manipulating the values extracted from these breakpoints.

3. Testing time-varying properties

```

1  # Get current time from last break point
2  current_time := time_api.c_timer.now_millis_last_bp;
3  # Check that it satisfies the 3 seconds limit
4  require current_time ≤ 3000;
```

¹ <https://github.com/ElNiak/PFV>

The time breakpoint is set in the PING frame event handler as presented in Listing 4:

4. Adding time breakpoint

```

1  object ping = {
2      around handle {
3          # [previous requirements]
4          ...
5          call time_api.c_timer.start; #add time break point
6          # [previous requirements]
7      }
8  }
```

This example demonstrates how simple it is to test a time-varying network property with a safety property thanks to the network-simulator assumption.

Three distinct implementations of the protocol were discerned. The first implementation consistently adhered to the specification by responding with PONG within the 3-second limit. In contrast, the second implementation displayed intermittent deviations from the desired behavior, indicating the necessity for further refinement. The third implementation consistently failed to meet the specification, exposing significant deficiencies.

Shadow’s capabilities were employed to introduce link jitter between the client and the server, simulating network conditions with varying packet delivery times. This additional element of uncertainty influenced the performance of the implementations. The previously flaky implementation, which occasionally deviated from the specification, now violated the time constraint more frequently.

The experiment’s determinism helped us to identify a specific seed value that leads to early connection failure in the flaky implementation.

Shadow’s debugging capabilities allow precise analysis of the two faulty implementations. By attaching a debugger to the implementations during testing, we can precisely identify which components were responsible for the deviations from the specified behavior.

4 Threat to validity for QUIC

This section provides an overview of how NSCT is applied to the QUIC protocol. It begins by defining the QUIC protocol and then discusses the modifications made to the formal model described in previous work [21, 46]. It also includes an analysis of the results obtained from the *picoquic* implementation, highlighting a specification violation. This issue was subsequently addressed through a pull request in the *picoquic* repository, fixing the error and aligning the implementation with the QUIC specifications.

QUIC is a modern transport protocol that combines the advantages of TCP (Transport Control Protocol) and TLS 1.3 (Transport Layer Security), while overcoming their limitations, as detailed in RFC9000 [31]. It introduces innovative secure communication methods at the transport layer. The RFC describes

how data are organized into frames and packets to ensure effective data segmentation, reliability, and control.

Tested implementation: *picoquic* is a research implementation of QUIC [15]. This implementation participated in the development of a QUIC standard by providing feedback. *picoquic* is written in C and consists of 103k lines of code. The tool incorporates various QUIC extensions and is currently under active development, making it an ideal choice for testing purposes. Moreover, QUIC has recently been chosen as the basis of HTTP/3 and is expected to handle a substantial portion of internet traffic in the coming years [46].

Table 1 summarizes the contribution to QUIC formal specification and the problems we found per RFC while testing *picoquic*:

	A. RFC9000	B. RFC9002	C. Ack Frequency
Previous works	Partially complete	/	/
Contributions	- Ack-delay - Idle timeout	- Congestion control (rtt calculation) - Loss recovery	90% of the draft
Problems found	Max retransmission	/	Misinterpretation in a frame field

Table 1: Summary of contributions to Ivy model and problems found in *picoquic*

A. Analysis of RFC9000: Our approach, integrating the time module and Shadow, has enabled enhancements to the existing QUIC model by incorporating time-related requirements as per RFC9000 specifications.

We focused on the idle timeout connection termination behavior of QUIC. QUIC outlines three primary methods for connection termination: immediate close, stateless reset, and idle timeout. We designed our tests to validate the implementation of these methods, particularly the idle timeout.

According to RFC9000 section 10.1, an endpoint restarts its idle timer upon receiving or sending ack-eliciting packets (i.e., the packet triggering ACK mechanism), ensuring that connections remain open during active communication.

To prevent overly brief idle timeouts, QUIC mandates an idle timeout period be at least three times the current Probe Timeout (PTO). This extension allows multiple opportunities for packet transmission before a timeout.

The connection in QUIC is automatically and *silently* closed, discarding its states if it remains idle beyond the minimum duration set by the `max_idle_timeout` transport parameter.

Our experiments revealed some discrepancies in the implementation of the idle timeout feature not in line with the standard behavior dictated by RFC9000.

We noticed deviations in the handling of retransmission thresholds and idle timeouts.

This test involved suspending packet transmission after a random period and observing if the connection closes silently, according to the specifications. However, our experiments revealed a deviation in the picoquic implementation. Rather than closing the connection after the idle timer expired, picoquic terminated it prematurely upon reaching a retransmission threshold. This behavior, probably influenced by TCP retransmission mechanism, deviates from RFC9000 standards, which require explicit notification through `CONNECTION_CLOSE` or `APPLICATION_CLOSE` frames for such terminations.

The discovered issue has been resolved through a pull request that was merged into the picoquic repository. This confirms the effectiveness of NSCT in detecting real-world anomalies in protocols.

B. Analysis of RFC9002 [58] This RFC discusses loss detection and congestion control in QUIC, differentiating it from TCP. It includes 37 mandatory specifications and 27 recommendations as per RFC2119. Key concepts introduced include the "probe timeout" (PTO) for managing congestion windows and the round-trip time (RTT) estimation process, comprising metrics like "min_rtt", "smoothed_rtt", and "rttvar". The RFC also details a sender-side congestion control mechanism, akin to TCP/NewReno, focusing on packet losses and Explicit Congestion Notification (ECN) [5, 24].

Our analysis involved implementing and testing the specified requirements and behaviors of RFC9002 in the context of congestion control and loss recovery, excluding the ECN component because it requires kernel support, which is not currently supported by Shadow. Future work could explore additional congestion control mechanisms like CUBIC [55] or BBR [14], and extensions such as QUIC-FEC [48]. Tests were conducted to evaluate the model's behavior under various network conditions, including loss, delay, and jitter, ensuring adherence to the RFC's guidelines.

While our formal specification of the RFC9002 did not identify specific problems in the *picoquic* implementation, it significantly contributed to refining the formal specification of QUIC, making it more precise and closely aligned with real-world scenarios.

C. Analysis of "QUIC Acknowledgement Frequency" [30] extension Currently in its draft-05 version, this extension enhances QUIC by allowing for delayed packet acknowledgments. It introduces the `min_ack_delay` transport parameter and two new frames: `ACK_FREQUENCY` and `IMMEDIATE_ACK`. The `ACK_FREQUENCY` frame adjusts acknowledgment rates based on the network state, while the `IMMEDIATE_ACK` frame assists connection liveliness.

In our examination, we analyzed the integration of this extension into the QUIC formal specification, focusing on the implications of delayed acknowledgments. During this process, an error was identified in the picoquic implementation; it incorrectly returned a `FRAME_ENCODING_ERROR` when processing an `ACK_FREQUENCY` frame. This issue, initially suspected to be a draft inconsistency,

was actually due to a misinterpretation of the "ACK-Eliciting Threshold" field in *picoquic*.

5 Discussion and future work

It is clear that the NSCT methodology, which involves various tools, is successful in uncovering new behaviors in protocol implementations, especially in terms of their temporal dynamics. However, combining multiple tools also brings about the intrinsic limitations of each tool and difficulties from their joint use. For example, employing network simulators, such as Shadow, comes with specific constraints. The necessity for simulators to depend on system calls for interacting with implementations restricts the range of implementations that can be tested. Moreover, busy loops (an anti-pattern not very used) reduce precision in time as they do not wait through system calls. In addition, the topology part of the NSCT testing process requires scenario-based simulations, where the behavior of the network within the simulated environment is predetermined.

A natural strategy to tackle these limitations is to replace certain tools used in the existing NSCT framework. For example, employing a network simulator that accommodates dynamic topology might alleviate the restrictions of scenario-based simulations, though it would necessitate modifications to existing automation scripts. Furthermore, substituting Ivy with any tool that implements the NCT methodology could be feasible, but this would require adapting the Ivy models.

In addition to improvements related to the joint use of several tools. Other future avenues may also be investigated. For example, a further improvement of the time module can allow for a more thorough verification of different types of properties. Closer integration between Ivy's generation process and the time module could make generating events at specific time points easier. This adaptation would enhance the tool's precision and overall usefulness.

Expanding the scope of the testing to include different congestion mechanisms in various protocols would provide more detailed insight into the details related to implementation and effectiveness.

Applying the methodology and tools discussed here to a broader range and scale of QUIC implementations would better validate and improve reliability and security.

Another promising area of research lies in the examination of the synergy between formal attacks models in network and Shadow's capabilities, especially considering the ongoing development of protocols like MPQUIC [23]. Currently, MPQUIC is still in its draft phase, grappling with significant security considerations between two main solutions. Given Shadow's unique ability to create custom network topologies, our methodology stands to offer substantial assistance. It enables the modeling of both solutions under consideration for MPQUIC, providing a comprehensive framework to assess their security implications and vulnerabilities.

Furthermore, another innovative approach is to leverage AI to simplify the creation of formal models from RFCs. This integration would greatly streamline the modeling process, making it more efficient and accessible. Additionally, a Graphical User Interface (GUI) for Ivy would allow users to engage in formal modeling without needing to understand the complexities of Ivy code, thereby making the process more user-friendly and approachable.

6 Conclusion

Protocol validation methods are diverse. Notably, the *INET* suite within the *OMNeT++* simulation library offers a powerful tool for network protocol validation [60]. Previous research used *INET* to simulate a QUIC model within *OMNeT++* [62, 63], but the evaluation was limited to protocol models and did not include real-world implementations. Studies like [2] used formal verification methods and simulations to validate complex protocol properties. This work measured the impact of specific attacks like Denial-of-Service on network parameters like energy consumption and computational effort, but did not use formal methods alongside simulations for verification.

Other research [7, 8, 40] attempted to enable the expression of time-varying network properties in the ISO standard to specify OSI protocols, known as Language Of Temporal Ordering Specification (*LOTOS*) [6]. While extensions to *LOTOS* offered varied expressiveness [8], subsequent Model-Based Testing (MBT) tools like *TorXakis* [59] only ensure guarantees about the tester’s side due to its lack of a network simulator. This limits its ability to verify time-varying network properties.

In [41], the authors compose their model with a network model to control the non-determinism of the network. This approach increases determinism, but it is not as powerful as NSCT, which extends determinism to IUT. [4] proposed using a test oracle on IUT traces, which allows offline verification of time-sensitive network properties that change based on network conditions. This approach avoids the high computational cost associated with online verification. However, this approach does not allow for the expression of time-dependent scenarios. Additionally, verifying traces does not facilitate the reproducibility of errors.

In this study, we propose an extension to Network-centric Compositional Testing (NCT). NCT is a simulation-based formal verification approach previously employed to validate QUIC implementations with Ivy. However, it has certain drawbacks; Ivy and NCT cannot capture time-varying network requirements or replicate experiments due to the inherent randomness of the methodology and the network. In addition, the extensive computational time required to scrutinize actual implementations of Internet protocols may affect protocol behaviors.

This study has successfully demonstrated the efficacy of *Network Simulator-centric Compositional Testing* (NSCT) in enhancing the verification of network protocols, particularly in addressing key challenges of NCT. NSCT, through the integration of the Ivy tool and the Shadow network simulator, effectively solves

several issues. These include the addition of time-varying network property verification, ensuring deterministic outcomes in protocol testing, and enhancing the reproducibility of test results. Our paper demonstrates the described method using a custom minimalist MiniP protocol. The application of NSCT in the picoquic implementation of QUIC identified a compliance error with time-varying network specifications, which was then rectified. This underscores the methodology's capability in managing complex, real-world network scenarios.

Additionally, a formal model is developed for RFC9002 that integrates congestion control and loss recovery into the existing QUIC model. The formal model of the "Acknowledgement Frequency" QUIC extension is also included.

Acknowledgement We would like to thank Maxime Piroux for his help to validate QUIC experiment results.

Artefacts The artefacts of this paper are available at <https://zenodo.org/doi/10.5281/zenodo.10819552>.

References

1. <https://groups.google.com/g/ns-3-users/c/NyX71jXHgr4?pli=1>, (Accessed on 12/10/2023)
2. Bernardeschi, C., Dini, G., Palmieri, M., Racciatti, F.: A framework for formal analysis and simulative evaluation of security attacks in wireless sensor networks. *Journal of Computer Virology and Hacking Techniques* **17**(3), 249–263 (Aug 2021). <https://doi.org/10.1007/s11416-021-00392-0>, <https://doi.org/10.1007/s11416-021-00392-0>
3. Bhargavan, K., Fournet, C., Kohlweiss, M., Pironti, A., Strub, P.Y.: Implementing tls with verified cryptographic security. In: 2013 IEEE Symposium on Security and Privacy. pp. 445–459. IEEE (2013)
4. Bishop, S., Fairbairn, M., Mehnert, H., Norrish, M., Ridge, T., Sewell, P., Smith, M., Wansbrough, K.: Engineering with logic: Rigorous test-oracle specification and validation for tcp/ip and the sockets api. *Journal of the ACM (JACM)* **66**(1), 1–77 (2018)
5. Black, D.L.: Rfc 8311: Relaxing restrictions on explicit congestion notification (ecn) experimentation (Jan 2018), <https://datatracker.ietf.org/doc/html/rfc8311>
6. Bolognesi, T., Brinksma, E.: Introduction to the iso specification language lotos. *Computer Networks and ISDN systems* **14**(1), 25–59 (1987)
7. Bolognesi, T., Lucidi, F.: A timed full lotos with time/action tree semantics. In: *Theories and Experiences for Real-Time System Development*, pp. 205–237. World Scientific (1994)
8. Bolognesi, T., Lucidi, F., Trigila, S.: Converging towards a timed lotos standard. *Computer Standards & Interfaces* **16**(2), 87–118 (1994)
9. Boyer, B., Corre, K., Legay, A., Sedwards, S.: Plasma-lab: A flexible, distributable statistical model checking library. In: Joshi, K., Siegle, M., Stoelinga, M., D’Argenio, P.R. (eds.) *Quantitative Evaluation of Systems*. pp. 160–164. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)

10. Bozic, J., Marsso, L., Mateescu, R., Wotawa, F.: A formal tls handshake model in int. In: 3rd Workshop on Models for Formal Analysis of Real Systems and 6th International Workshop on Verification and Program Transformation, MARSVPT 2018. pp. 1–40 (2018)
11. Breslau, L., Estrin, D., Fall, K., Floyd, S., Heidemann, J., Helmy, A., Huang, P., McCanne, S., Varadhan, K., Xu, Y., et al.: Advances in network simulation. *Computer* **33**(5), 59–67 (2000)
12. Cadar, C., Dunbar, D., Engler, D.R., et al.: Klee: Unassisted and automatic generation of high-coverage tests for complex systems programs. In: OSDI. vol. 8, pp. 209–224 (2008)
13. Cappart, Q., Limbrée, C., Schaus, P., Quilbeuf, J., Traonouez, L.M., Legay, A.: Verification of interlocking systems using statistical model checking. In: 2017 IEEE 18th International Symposium on High Assurance Systems Engineering (HASE). pp. 61–68 (2017). <https://doi.org/10.1109/HASE.2017.10>
14. Cardwell, N., Cheng, Y., Yeganeh, S.H., Swett, I., Jacobson, V.: Bbr congestion control, <https://datatracker.ietf.org/doc/html/draft-cardwell-iccr-g-bbr-congestion-control>
15. Christian Huitema: picoquic, <https://github.com/private-octopus/picoquic>, 4f11445
16. Chudnov, A., Collins, N., Cook, B., Dodds, J., Huffman, B., MacCárthaigh, C., Magill, S., Mertens, E., Mullen, E., Tasiran, S., et al.: Continuous formal verification of amazon s2n. In: Computer Aided Verification: 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14–17, 2018, Proceedings, Part II 30. pp. 430–446. Springer (2018)
17. Clarke, E.M., Donzé, A., Legay, A.: On simulation-based probabilistic model checking of mixed-analog circuits. *Formal Methods Syst. Des.* **36**(2), 97–113 (2010). <https://doi.org/10.1007/S10703-009-0076-Y>, <https://doi.org/10.1007/s10703-009-0076-y>
18. Clarke, E.M., Emerson, E.A.: Design and synthesis of synchronization skeletons using branching time temporal logic. In: Workshop on logic of programs. pp. 52–71. Springer (1981)
19. Clarke, E.M., Grumberg, O., Kroening, D., Peled, D.A., Veith, H.: Model checking, 2nd Edition. MIT Press (2018), <https://mitpress.mit.edu/books/model-checking-second-edition>
20. Classen, A., Heymans, P., Schobbens, P.Y., Legay, A.: Symbolic model checking of software product lines. In: Proceedings of the 33rd International Conference on Software Engineering. pp. 321–330 (2011)
21. Crochet, C., Rousseaux, T., Piroux, M., Sambon, J.F., Legay, A.: Verifying quic implementations using ivy. Proceedings of the 2021 Workshop on Evolution, Performance and Interoperability of QUIC (2021). <https://doi.org/10.1145/3488660.3493803>
22. David, A., Larsen, K.G., Legay, A., Mikučionis, M., Poulsen, D.B.: Uppaal smc tutorial. *International journal on software tools for technology transfer* **17**, 397–415 (2015)
23. De Coninck, Q., Bonaventure, O.: Multipath QUIC. In: Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies. ACM (Nov 2017). <https://doi.org/10.1145/3143361.3143370>, <https://doi.org/10.1145/3143361.3143370>
24. Floyd, S., Ramakrishnan, D.K.K., Black, D.L.: Rfc 3168: The addition of explicit congestion notification (ecn) to ip (Sep 2001), <https://datatracker.ietf.org/doc/html/rfc3168>

25. Fujimoto, R.M.: Parallel and distributed simulation systems. In: Proceeding of the 2001 Winter Simulation Conference (Cat. No. 01CH37304). vol. 1, pp. 147–157. IEEE (2001)
26. Fujimoto, R.M., Riley, G.F., Perumalla, K.S.: Network Simulators. Springer International Publishing, Cham (2007). <https://doi.org/10.1007/978-3-031-79977-8>
27. Holzmann, G.J.: The model checker SPIN. IEEE Trans. Software Eng. **23**(5), 279–295 (1997). <https://doi.org/10.1109/32.588521>, <https://doi.org/10.1109/32.588521>
28. Issariyakul, T., Hossain, E., Issariyakul, T., Hossain, E.: Introduction to network simulator 2 (NS2). Springer (2009)
29. Iyengar, J., Swett, I.: Rfc 9002, <https://www.rfc-editor.org/rfc/rfc9002.html>
30. Iyengar, J., Swett, I., Kühlewind, M.: Quic acknowledgement frequency, <https://datatracker.ietf.org/doc/html/draft-ietf-quic-ack-frequency-05>
31. Iyengar, J., Thomson, M.: Rfc 9000, <https://www.rfc-editor.org/rfc/rfc9000>
32. Jansen, R., Hopper, N.J.: Shadow: Running tor in a box for accurate and efficient experimentation (2011)
33. Jansen, R., Newsome, J., Wails, R.: Co-opting linux processes for High-Performance network simulation. In: 2022 USENIX Annual Technical Conference (USENIX ATC 22). pp. 327–350. USENIX Association, Carlsbad, CA (Jul 2022), <https://www.usenix.org/conference/atc22/presentation/jansen>
34. Katoen, J.P.: The probabilistic model checking landscape. In: Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science. p. 31–45. LICS '16, Association for Computing Machinery, New York, NY, USA (2016). <https://doi.org/10.1145/2933575.2934574>, <https://doi.org/10.1145/2933575.2934574>
35. Kumar, R., Stoelinga, M.: Quantitative security and safety analysis with attack-fault trees. In: 2017 IEEE 18th International Symposium on High Assurance Systems Engineering (HASE). pp. 25–32 (2017). <https://doi.org/10.1109/HASE.2017.12>
36. Larsen, K.G., Legay, A.: 30 years of statistical model checking. In: Margaria, T., Steffen, B. (eds.) Leveraging Applications of Formal Methods, Verification and Validation: Verification Principles - 9th International Symposium on Leveraging Applications of Formal Methods, ISoLA 2020, Rhodes, Greece, October 20–30, 2020, Proceedings, Part I. Lecture Notes in Computer Science, vol. 12476, pp. 325–330. Springer (2020). https://doi.org/10.1007/978-3-030-61362-4_18, https://doi.org/10.1007/978-3-030-61362-4_18
37. Larsen, K.G., Mikucionis, M., Nielsen, B.: Uppaal tron user manual. CISS, BRICS, Aalborg University, Aalborg, Denmark (2009)
38. Lee, H., Seibert, J., Fistrovic, D., Killian, C., Nita-Rotaru, C.: Gatling: Automatic performance attack discovery in large-scale distributed systems. ACM Transactions on Information and System Security (TISSEC) **17**(4), 1–34 (2015)
39. Legay, A., Sedwards, S.: On statistical model checking with plasma. In: The 8th International Symposium on Theoretical Aspects of Software Engineering (2014)
40. Léonard, L., Leduc, G.: An introduction to et-lotos for the description of time-sensitive systems. Computer networks and ISDN systems **29**(3), 271–292 (1997)
41. Li, Y., Pierce, B.C., Zdancewic, S.: Model-based testing of networked applications. In: Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis. pp. 529–539 (2021)

42. Lounas, R., Jafri, N., Legay, A., Mezghiche, M., Lanet, J.L.: A formal verification of safe update point detection in dynamic software updating. In: *Risks and Security of Internet and Systems: 11th International Conference, CRIStIS 2016, Roscoff, France, September 5-7, 2016, Revised Selected Papers 11*. pp. 31–45. Springer (2017)
43. McMillan, K.L.: *Symbolic model checking*. Kluwer (1993). <https://doi.org/10.1007/978-1-4615-3190-6>, <https://doi.org/10.1007/978-1-4615-3190-6>
44. McMillan, K.L., Padon, O.: Ivy: A multi-modal verification tool for distributed algorithms. *Computer Aided Verification* p. 190–202 (2020). https://doi.org/10.1007/978-3-030-53291-8_12
45. McMillan, K.L., Zuck, L.D.: Compositional testing of internet protocols. *2019 IEEE Cybersecurity Development (SecDev)* (2019). <https://doi.org/10.1109/secdev.2019.00031>
46. McMillan, K.L., Zuck, L.D.: Formal specification and testing of quic. *Proceedings of the ACM Special Interest Group on Data Communication* (2019). <https://doi.org/10.1145/3341302.3342087>
47. Merkel, D.: Docker: lightweight linux containers for consistent development and deployment. *Linux journal* **2014**(239), 2 (2014)
48. Michel, F., De Coninck, Q., Bonaventure, O.: Quic-fec: Bringing the benefits of forward erasure correction to quic. In: *2019 IFIP Networking Conference (IFIP Networking)*. pp. 1–9 (2019). <https://doi.org/10.23919/IFIPNetworking.2019.8816838>
49. Ngo, V.C., Legay, A., Joloboff, V.: PSCV: A runtime verification tool for probabilistic systemc models. In: Chaudhuri, S., Farzan, A. (eds.) *Computer Aided Verification - 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part I. Lecture Notes in Computer Science*, vol. 9779, pp. 84–91. Springer (2016). https://doi.org/10.1007/978-3-319-41528-4_5, https://doi.org/10.1007/978-3-319-41528-4_5
50. Offutt, J., Abdurazik, A.: Generating tests from uml specifications. In: *International Conference on the Unified Modeling Language*. pp. 416–429. Springer (1999)
51. Padon, O., McMillan, K.L., Panda, A., Sagiv, M., Shoham, S.: Ivy: Safety verification by interactive generalization. *ACM SIGPLAN Notices* **51**(6), 614–630 (2016). <https://doi.org/10.1145/2980983.2908118>
52. Paris, J., Arts, T.: Automatic testing of tcp/ip implementations using quickcheck. In: *Proceedings of the 8th ACM SIGPLAN Workshop on Erlang*. pp. 83–92 (2009)
53. Pnueli, A.: The temporal logic of programs. In: *18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October - 1 November 1977*. pp. 46–57. IEEE Computer Society (1977). <https://doi.org/10.1109/SFCS.1977.32>, <https://doi.org/10.1109/SFCS.1977.32>
54. Rath, F., Schemmel, D., Wehrle, K.: Interoperability-guided testing of quic implementations using symbolic execution. In: *Proceedings of the Workshop on the Evolution, Performance, and Interoperability of QUIC*. pp. 15–21 (2018)
55. Rhee, I., Xu, L., Ha, S., Zimmermann, A., Eggert, L., Scheffenegger, R.: Rfc 8312: Cubic for fast long-distance networks (Feb 2018), <https://datatracker.ietf.org/doc/html/rfc8312>
56. Riley, G.F., Henderson, T.R.: The ns-3 network simulator. In: *Modeling and tools for network simulation*, pp. 15–34. Springer (2010)
57. Tazaki, H., Uarbani, F., Mancini, E., Lacage, M., Camara, D., Turletti, T., Dabous, W.: Direct code execution: Revisiting library os architecture for reproducible network experiments. In: *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies*. pp. 217–228 (2013)

58. Thomson, M., Turner, S.: Rfc 9001, <https://www.rfc-editor.org/rfc/rfc9001.html>
59. Tretmans, G., van de Laar, P.: Model-based testing with torxakis: the mysteries of dropbox revisited (2019)
60. Varga, A.: Omnet++. In: Modeling and tools for network simulation, pp. 35–59. Springer (2010)
61. Veanes, M., Campbell, C., Grieskamp, W., Schulte, W., Tillmann, N., Nachmanson, L.: Model-based testing of object-oriented reactive systems with spec explorer. Formal Methods and Testing: An Outcome of the FORTEST Network, Revised Selected Papers pp. 39–76 (2008)
62. Völker, T., Volodina, E., Tüxen, M., Rathgeb, E.P.: A quic simulation model for inet and its application to the acknowledgment ratio issue. In: 2020 IFIP Networking Conference (Networking). pp. 737–742. IEEE (2020)
63. Volodina, E., Rathgeb, E.P.: Impact of ack scaling policies on quic performance. In: 2021 IEEE 46th Conference on Local Computer Networks (LCN). pp. 41–48 (2021). <https://doi.org/10.1109/LCN52139.2021.9524947>