

# Simulation-based Evaluation of a Synchronous Transaction Model for Time-Sensitive Software-Defined Networks

Tobias Haugg, Mohammad Fazel Soltani, Timo Häckel, Philipp Meyer, Franz Korf, and Thomas C. Schmidt

*Dept. Computer Science, Hamburg University of Applied Sciences, Germany*  
{tobias.haugg, mohammadfazel.soltani, timo.haeckel, philipp.meyer,  
franz.korf, t.schmidt}@haw-hamburg.de

## Abstract

Real-time networks based on Ethernet require robust quality-of-service for time-critical traffic. The Time-Sensitive Networking (TSN) collection of standards enables this in real-time environments like vehicle on-board networks. Runtime reconfigurations in TSN must respect the deadlines of real-time traffic. Software-Defined Networking (SDN) moves the control plane of network devices to the SDN controller, making these networks programmable. This allows reconfigurations from a central point in the network. In this work, we present a transactional model for network reconfigurations that are synchronously executed in all network devices. We evaluate its performance in a case study against non-transactional reconfigurations and show that synchronous transactions enable consistency for reconfigurations in TSN without increased latencies for real-time frames.

## 1 Introduction

In recent years, Ethernet is playing an increasingly important role, even finding its way into real-time applications. Complementary protocols such as IEEE 802.1Q [6] Time-Sensitive Networking (TSN) have proven to meet the real-time and robustness requirements of these environments. On the other hand, Software-Defined Networking (SDN) [9] has revolutionized campus and data center networks. Separating the control from the data plane of network devices at a central control unit enables simple and fast-forwarding at devices, while high level control applications with global network knowledge can steer the entire network.

One challenge in frequently changing networks is reconfiguration at runtime, which is often necessary when flows are added or removed sporadically, and down time is not an option. The biggest problem faced is maintaining a consistent state without affecting network performance [1], which is especially important in real-time environments. Transactions can solve these problems in distributed systems by providing the *ACID* properties *atomicity*, *consistency*, *isolation*, and *durability* [11]. This way transactions guarantee that the network is always in a consistent state no matter whether a reconfiguration succeeds or not. In real-time environments the synchronous execution of transactions over multiple network devices also plays an important role when configuring traffic with low-latency requirements and no tolerance to packet loss.

In this work, we present our time synchronous transaction model and compare it with a non-transactional reconfiguration approach, both of which we implement using NETCONF and study their performance. To evaluate the different implementations, we use the discrete event simulation platform OMNeT++. Our implementations extend our previously published simulation framework *Software-Defined Networking for Communication over Real-Time Ethernet (SDN4CoRE)* [4], which enables the simulation of programmable (software-defined) real-time Ether-networks.

The paper is structured as follows. Section 2 provides background knowledge and related work. Section 3 shows our concept of synchronous transactions. Section 4 compares the performance of synchronous transactions with non-transactional network reconfiguration in a case study. Finally, section 5 draws a conclusion and gives an outlook on future work.

## 2 Background & Related Work

Our time-synchronous transaction model enables reconfiguration of networks combining TSN and SDN. This section introduces into TSN, SDN, Transactions, and related work.

### 2.1 Time-Sensitive Networking

The TSN real-time Ethernet protocol is a collection of standards specifically designed for use in industrial control equipment and communication networks within vehicles. TSN is defined in IEEE 802.1Q [6] and bases on Ethernet frames with Q-Tag that have a VLAN ID and priority code point (from 0 - low to 7 - high). To enable heterogenous Quality-of-Service (QoS) guarantees for each individual priority each port has a frame selector instance that decides when which packet is sent.

Figure 1 shows the frame selection. Frames can be assigned to one of eight priorities. For each priority, the frame selector has a queue, a Transmission Selection Algorithm, and a Transmission Gate. If more than one path is open to the Transmission Selection, the frame is taken from the queue with the highest priority and sent to the interface.

Transmission Selection Algorithm instances can implement different frame selection algorithms.

An example of this is the Credit Based Shaping (CBS) algorithm, which ensures compliance with a reserved bandwidth. If the Transmission Selection Algorithm allows a frame to be sent, the next thing to consider is the transmission gate.

The gates have one of two possible states “OPEN” or “CLOSED”. Sending a frame is only allowed when the “OPEN” state is active. The state change can be timed based on a configuration called Gate Control List (GCL). This enables scheduling of Time Division Multiple Access (TDMA) traffic by opening individual Transmission Gates in specific time intervals. The accuracy of such a schedule depends on which devices are involved in the scheduling. In a fully scheduled network, all endpoints and switches implement a common TDMA schedule.

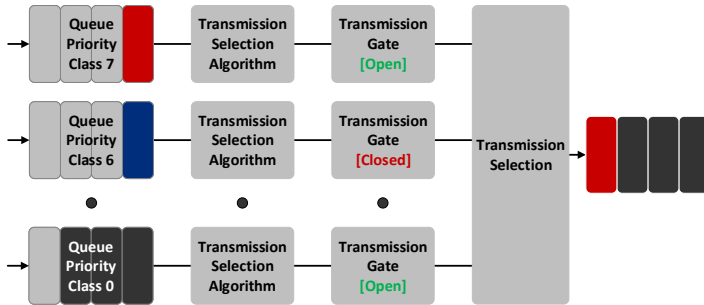


Figure 1: TSN scheduling and forwarding [6]

## 2.2 Software-Defined Networking

Software-Defined Networking (SDN) [9] separates the control logic (control plane) from the underlying switches that forward the traffic (data plane) [8]. Network devices become programmable by a central SDN controller using open standard protocols such as OpenFlow. OpenFlow switches forward incoming packets based on a programmable flow table. A flow entry matches a subset of Layer 2 to Layer 4 header fields and contains actions, such as discard or forward, that are performed when traffic matches. SDN controller applications implement the behavior of the network, e.g. routing protocols, and populate the flow tables in the switches.

SDN can reduce the complexity and increase the adaptability of real-time Ether-networks to cope with the growing dynamics introduced, for example, by Service-Oriented Architecture (SOA). In previous work, we integrated TSN with SDN [5] to control dynamic bandwidth reservation with a central SDN controller without delay penalty for real-time communication. Dynamic control of scheduled traffic is more complex because it must be coordinated between multiple senders across multiple hops in the network. In this paper, we explore different solutions for network schedule reconfiguration.

## 2.3 Transactional Network (Re-)Configuration

Transactions are commonly used to coordinate changes in distributed systems [11]. The ACID properties can guarantee the validity of a transaction despite the occurrence of errors:

- *Atomicity*: A transaction composed of multiple operations either succeeds completely, or fails completely. If a transaction fails the system is left unchanged.
- *Consistency*: A transaction always transfers the system from one valid state to another.
- *Isolation*: Concurrent transactions behave exactly as if they were executed sequentially.
- *Durability*: When a transaction has been committed it remains committed until the next changes are made to the system.

The properties of ACID transactions can be used in SDN environments to maintain the consistency of the network, e.g. by extending the OpenFlow protocol [2]. Cui et al. show how transactions can bundle changes to multiple flow rules of a network function [1]. OpenFlow is well suited to update flows in the network. It does not support TSN GCL configuration.

The NETCONF protocol (RFC 6241 [3]) can be used for reconfiguration of TSN modules and the OpenFlow table at the same time via existing YANG data models. NETCONF provides useful features to implement transactions. Each device has multiple configuration data stores. Each data store can be locked by exactly one client guaranteeing isolation until the store is unlocked. One data store contains the active device configuration (*running*) and changes to it are directly applied and activated. *Candidate* datastores contain a copy of the running configuration and changes are applied only to the copy so that the full configuration can be verified and later committed and thus applied to the running datastore.

In this work, we extend the NETCONF transaction mechanism to allow synchronously executed transactions over multiple network devices.

## 3 Synchronous Transactions for Network Reconfiguration

Our time-synchronous transaction model is coordinated by an SDN controller that reconfigures a set of switches using the NETCONF protocol.

Figure 2 shows the simulation environment used for our implementation, which is based on the discrete event simulator OMNeT++<sup>1</sup> (version 5.5.1). The INET framework<sup>2</sup> (version 3.x) implements internet and transport protocols, our fork of the OpenFlowOmnetSuite[7] enables SDN with the OpenFlow protocol, our CoRE4INET framework [10] implements TSN. We incorporated the transaction model into the SDN4CoRE framework, which we presented in previous work [4]. It adds a programming option for real-time Ether-networks and provides an implementation of the NETCONF protocol. Our frameworks are published and available as open source<sup>3</sup>.

A few assumptions were made to simplify the implementation: A reliable transport protocol is used, there are no hardware or network errors, and each switch supports the same TSN and SDN features. The clocks of all devices in the network are synchronized with a maximum inaccuracy of  $\pm 500$  ns.




Figure 2: Simulation environment

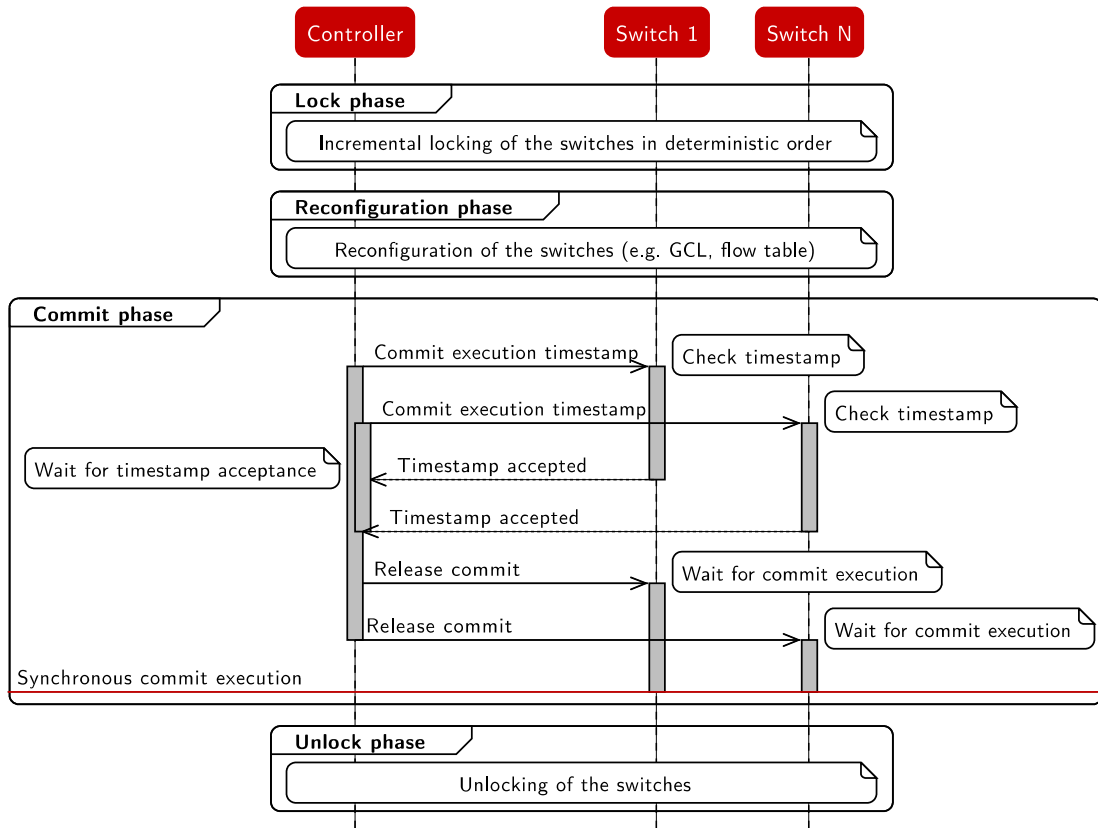


Figure 3: Simplified sequence diagram of the four phases in a successful transaction.

<sup>1</sup> [omnetpp.org](http://omnetpp.org) <sup>2</sup> [inet.omnetpp.org](http://inet.omnetpp.org) <sup>3</sup> [sim.core-rg.de](http://sim.core-rg.de) and [github.com/CoRE-RG](https://github.com/CoRE-RG)

Our synchronous transaction model consists of four phases: *Lock*, *reconfiguration*, *commit*, and *unlock*. Figure 3 depicts a successful transaction in a sequence diagram. The lock, reconfiguration, and unlock phases use standard NETCONF operations and are thus shown in simplified form. The commit phase is shown in detail. It implements the two phase commit protocol [11] and adds support for synchronized reconfiguration.

At first, the lock phase reserves all resources needed for the transaction. The controller attempts to acquire a lock for each switch whose configuration is supposed to be modified. To avoid deadlocks, the switches are contacted one after the other in the order of their MAC addresses. If a lock cannot be acquired, all previously obtained locks are released, and the transaction terminates unsuccessfully.

The second phase is the reconfiguration phase. For each switch, a candidate configuration is instantiated that is a copy of the previously locked running configuration. The candidate configuration is locked immediately after it is instantiated and is used to apply any reconfigurations. If an error occurs during the copy operation, while locking the candidate configuration, or during any of the reconfigurations, all the changes made are undone and the running configuration remains unchanged.

In the third phase, called commit phase, the candidate configuration of each switch becomes the new running configuration. A network-wide commit execution time is used to synchronously activate the new configurations in all switches. The controller calculates the commit execution time and performs a two-stage commit, first checking that each switch is ready and accepts the timestamp, and then releasing the commit for execution at the agreed time. The timestamp for the commit execution is calculated based on the Worst-Case Execution Time (WCET) of the commit phase. For this, the accumulated maximum latencies for all messages and processing times at the controller and switches are summed up. The WCET is independent from the complexity of the reconfigurations since they have already been applied to the candidate configuration in the previous phase. This time span can then be used to either directly determine a commit execution timestamp or the next possible start of a network period for the commit execution. All devices apply updates in the candidate configuration to the running configuration when the synchronous timestamp is reached.


At the end of the transaction, the fourth phase unlocks all resources and cleans up after the transaction. The unlock phase is always executed, regardless of whether any of the previously described phases failed. All unconfirmed changes and the old configuration are deleted. At last, all previously acquired locks are released.

## 4 Case Study on Network Reconfiguration Consistency

We compare the performance of the transactional model with a non-transactional approach using the impact of reconfigurations on the end-to-end latency of frames from real-time flows.

Figure 4 shows our evaluation network, which consists of four nodes connected via two Time-Sensitive Software-Defined Networking (TSSDN) switches. At the control plane, these switches are connected to a Time-Sensitive Software-Defined Networking (TSSDN) controller via two Ethernet switches that introduce a separate delay between controller and Time-Sensitive Software-Defined Networking (TSSDN) switches. Each link is 10 meters long and has a speed of 100 Mbit/s. Also, each switch has a forwarding delay of 3  $\mu$ s.

Node1, node2, and node3 send one full-size Ethernet packet every millisecond over their respective real-time flows 1, 2, and 3, which are received by node4. On all devices a TDMA schedule is used repeating every 1 ms. All packets from flow 1, 2 and 3 share the same priority, and therefore the same time slot in the schedule. The time slot always has to adapt to the



**Figure 5:** End-to-end latency for non-synchronous, non-transactional network updates

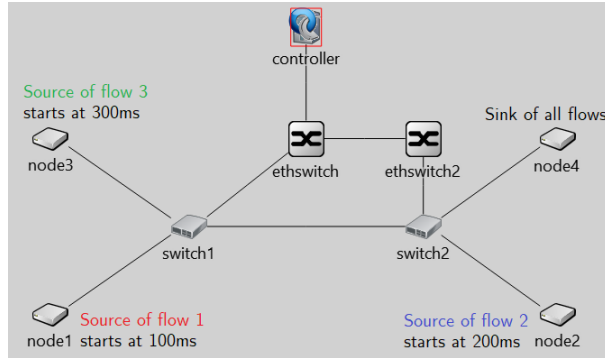
number of active real time flows. Switch1 has to offer a time slot long enough to forward the packets from flow 1 and 3 and switch2 to has to offer enough time to forward flow 1, 2 and 3 during one period. To avoid interference from cross traffic, a guard band is added to close all gates and guarantee that the band is clear when real-time traffic arrives.

We calculate the expected latencies of the real-time flows based on the transmission delay over the links, and the forwarding delay in the switches that a packet must pass. This results in an expected latency of  $373.35\ \mu\text{s}$  for packets originating from node1 and node3, and  $247.9\ \mu\text{s}$  for packets from node2.


Traffic at the nodes is started sequentially with an offset of 100 ms to observe the effects of the reconfigurations on existing and added traffic. After a node is activated, two reconfigurations are performed on the switches. First, the flow table is modified to forward the new packets correctly; second, the GCL schedule is updated to accommodate an additional packet of the new real-time flow in the time slot for scheduled traffic. A transaction bundles both modifications per activated node, while in the non-transactional approach these modifications are performed independently. In total, we have six reconfigurations (U1-U6) using the non-transactional approach and three transactions (T1-T3) combining related modifications.

Figure 5 shows the end-to-end latency for the non-transactional reconfiguration. The latency increases drastically during the third and fourth reconfigurations (U3 and U4) and then remains constant until the next flow is added. During U5 and U6, the latency of flow 1 increases again and then remains constant at a fairly high level together with flow 3. The latency of flow 2 decreases but is still above the previously calculated latency of  $247.9\ \mu\text{s}$ .


Figure 6 gives a more detailed view of reconfigurations U3 and U4 between 200 ms and 210 ms simulation time. After reconfiguration U3 adds a flow rule for flow 2, the flow is forwarded



**Figure 4:** Network topology of the case study



**Figure 6:** Detailed view of the end-to-end latency during reconfigurations U3 and U4 with non-synchronous, non-transactional network updates



**Figure 7:** End-to-end latency for time-synchronous, transactional network updates

through switch2. The latency for flow 1 spikes because the first packet from flow 2 takes away the time slot originally reserved for flow 1 and the GCL schedule has not yet been changed by reconfiguration U4. Thus, the packet from flow 1 misses its time slot and must wait one period before it can be transmitted. In the next period, the packet from flow 1 is transmitted, while the newly arrived packet from flow 2 must now wait for one period. Reconfiguration U4 updates the GCL to now accommodate two full-size packets per period, one for flow 1 and one for flow 2. This results in constant but severely delayed latencies. Similar behavior can be observed after U5 and U6. The drastically increased latencies of flow 1 and 3 show that the time-guarantees of the TSN network are no longer met, making this non-transactional approach inapplicable in real-time networks.

Figure 7 shows the results for the time-synchronous transactions. In the transactional approach, the two reconfigurations that need to be performed for each new node are combined into one transaction. In addition, the transaction is executed synchronously on both switches, reducing the period of inconsistency in the network to the accuracy of time synchronization. With time-synchronous transactions, the latency of all real-time flows stays constant and matches the calculated latencies throughout the simulation. The difference in latencies compared to the non-transactional approach is caused by the way the reconfigurations are bundled and executed.

Comparing the results of Figure 7 with the non-transactional approach (Figure 5), we can see that they differ not only in the end-to-end latency of the traffic, but also in the execution time of reconfigurations. Although the reconfiguration starts at the same time in both approaches, we can observe a delay in the execution of the transactions. This delay is caused by the synchronous commit execution based on the WCET.

Our results show that the synchronous transaction model is indeed able to perform network reconfiguration for real-time traffic without delay penalties for individual frames, making it possible to reconfigure TSN networks during runtime.

## 5 Conclusion & Outlook

In this paper, we presented our approach for synchronous transactional runtime reconfigurations in TSN. Our case study showed, due to inconsistencies in reconfigurations the non-transactional approach led to drastically increased latencies and has been shown to be unsuitable for time critical traffic. In contrast, the time-synchronous transactional model was able to reconfigure the TSN devices during runtime without affecting the latency of existing and added real-time traffic. By merging related reconfigurations into one transaction that is executed synchronously in all network devices, the network can be seamlessly transition between consistent states.

Future work will investigate the performance of synchronous transactions in TSN under various forms of modifications such as deleting and rerouting existing real-time flows. Furthermore, we will compare several forms of commit synchronization, e.g., commit execution in the middle of a period as opposed to the beginning of the period might be of interest.

## Acknowledgments

This work is funded by the Federal Ministry of Education and Research of Germany (BMBF) within the SecVI project.

## References

- [1] J. Cui, S. Zhou, H. Zhong, Y. Xu, and K. Sha. Transaction-Based Flow Rule Conflict Detection and Resolution in SDN. In *2018 27th International Conference on Computer Communication and Networks (ICCCN)*, pages 1–9, July 2018.
- [2] M. Curic, Z. Despotovic, A. Hecker, and G. Carle. Transactional Network Updates in SDN. In *2018 European Conference on Networks and Communications (EuCNC)*, pages 203–208, 2018.
- [3] R. Enns, M. Bjorklund, J. Schoenwaelder, and A. Bierman. Network Configuration Protocol (NETCONF). RFC 6241, IETF, June 2011.
- [4] Timo Häckel, Philipp Meyer, Franz Korf, and Thomas C. Schmidt. SDN4CoRE: A Simulation Model for Software-Defined Networking for Communication over Real-Time Ethernet. In Meyo Zongo, Antonio Virdis, Vladimir Vesely, Zeynep Vatandas, Asanga Udugama, Koojana Kuladinithi, Michael Kirsche, and Anna Förster, editors, *Proc. of the 6th Int. OMNeT++ Community Summit 2019*, volume 66 of *EPiC Series in Computing*, pages 24–31. EasyChair, December 2019.
- [5] Timo Häckel, Philipp Meyer, Franz Korf, and Thomas C. Schmidt. Software-Defined Networks Supporting Time-Sensitive In-Vehicular Communication. In *2019 IEEE 89th Vehicular Technology Conference (VTC2019-Spring)*, pages 1–5, Piscataway, NJ, USA, April 2019. IEEE Press.
- [6] IEEE 802.1 Working Group. IEEE Standard for Local and Metropolitan Area Network–Bridges and Bridged Networks. Std. 802.1Q-2018 (Revision of IEEE Std 802.1Q-2014), IEEE, July 2018.



- [7] Dominik Klein and Michael Jarschel. An OpenFlow extension for the OMNeT++ INET framework. In *Proceedings of the 6th International ICST Conference on Simulation Tools and Techniques*, SimuTools '13, pages 322–329, Brussels, BEL, 2013. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).
- [8] D. Kreutz, F. M. V. Ramos, P. E. Veríssimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig. Software-Defined Networking: A Comprehensive Survey. *Proceedings of the IEEE*, 103(1):14–76, January 2015.
- [9] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. OpenFlow: Enabling Innovation in Campus Networks. *ACM SIGCOMM Computer Communication Review*, 38(2):69–74, 2008.
- [10] Philipp Meyer, Franz Korf, Till Steinbach, and Thomas C Schmidt. Simulation of Mixed Critical In-vehicular Networks. In *Recent Advances in Network Simulation*, pages 317–345. Springer, 2019.
- [11] Maarten van Steen and Andrew S. Tanenbaum. *Distributed systems*. Pearson Education, 2017.