



Hochschule für Angewandte  
Wissenschaften Hamburg  
*Hamburg University of Applied Sciences*



Technische Universität München

Freie Universität



Berlin

# Towards an Ecosystem for Reproducible Research in Computer Networking

Quirin Scheitle (TUM), Matthias Wählisch (FU  
Berlin), Oliver Gasser (TUM), Thomas C. Schmidt  
(HAW Hamburg), Georg Carle (TUM)

# Recap: What do we want?




Figure 1: Motivating experiments: Hadoop traffic interferes with (a) PTPd, (b) memcached and (c) Naiad traffic.

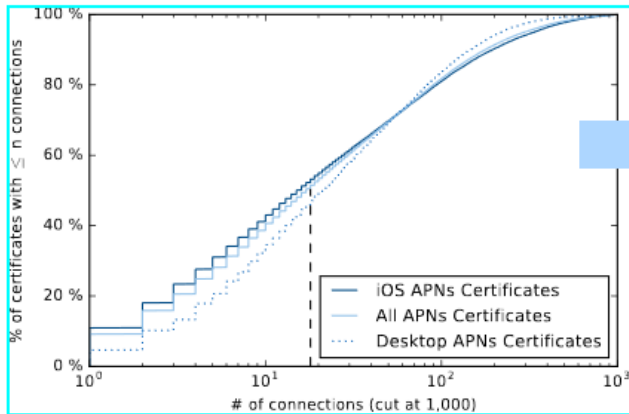
Setup	50 <sup>th</sup> %	99 <sup>th</sup> %
one host, idle network	85	126µs
two hosts, shared switch	110	130µs
shared source host, shared egress port	228	268µs
shared dest. host, shared ingress port	125	278µs
shared host, shared ingress and egress	221	229µs
two hosts, shared switch queue	1,920	2,100µs

Table 1: Median and 99<sup>th</sup> percentile latencies observed as ping and Iperf share various parts of the network.

in §6) and measure the effects.

**1. Clock Synchronization** Precise clock synchronization is important to distributed systems such as Google's Spanner [11]. PTPd offers microsecond-granularity time synchronization from a time server to machines on a local network. In Figure 1a, we show a timeline of PTPd synchronizing a host clock on both an idle network and when sharing the network with Hadoop. In the shared case, Hadoop's shuffle phases causes queuing which delays PTPd's synchronization marks. This

[Grosvenor et al., NSDI'15]




(a) ~50% of certificates observed with more than 17 connections, 9% of certificates only observed once (5% desktop, 11% iOS).

[Wachs et al., TMA'17]

Figure 1c / 3c

Figure 1c (Page 2) is used as a motivational experiment to show that Hadoop MapReduce is capable of interfering with the behaviour of Naiad. These results are repeated in Figure 3c (Page 7) with the addition of QJump to show that QJump is capable of resolving this interference.



### Software Required

- In addition to these general software requirements, we also used:
- QJump kernel module
  - QJump application utility
  - WDR head 5.2.3 (using the patched example binary test)
  - Hadoop MapReduce 4.5.0 (installed from Cloudera packages, instructions here)

Code Issues Pull requests Projects Wiki Insights

Branch: master cca-privacy / analyses / results / 1\_connections\_per\_certificate.ipynb Find file Copy path

quirins initial commit 3a53ce9 on 2 May

1 contributor

176 lines (175 sloc) | 34.8 KB Raw Blame History

### Analysis of number of connections per client certificate

```
In [3]: import numpy as np
import matplotlib.pyplot as plt
from matplotlib.ticker import FormatStrFormatter
import math
matplotlib inline

In [4]: no_v_rttc = np.loadtxt("1_connections_per_certificate.sh.csv.countonly", delimiter=",")
no_v_rttc_sort = np.sort(no_v_rttc)

ios = np.loadtxt("1_connections_per_certificate_ios.sh.csv.countonly", delimiter=",")
ios_sort = np.sort(ios)
macos = np.loadtxt("1_connections_per_certificate_macos.sh.csv.countonly", delimiter=",")
macos_sort = np.sort(macos)

# drop excessive connections from build cluster
no_v_rttc_sort = no_v_rttc_sort[no_v_rttc_sort < 1000]
ios_sort = ios_sort[ios_sort < 1000]
macos_sort = macos_sort[macos_sort < 1000]

Plot definitions

In [5]: ylabel = "% of certificates with ≤ n connections"
xlabel = "# of connections (cut at 1,000)"
xmin = 0
ymin = 0
path = ".,/"
```

# Recap: Where are we?



## Workshop on Models, Methods and Tools for Reproducible Network Research (MoMeTools)

Karlsruhe, Germany, August 25, 2003  
In conjunction with **ACM SIGCOMM 2003**

- Main
- Formatting/Submission
- Call for Papers
- Workshop program
- ACM Digital Library Proceedings

### Workshop on Models, Methods and Tools for Reproducible Network Research (MoMeTools)

Compared with other scientific areas such as experimental physics, network research appears significantly less mature concerning methodology. The goal of this workshop is to critically assess the current models, methods and tools of network research for identifying shortcomings of the state-of-the-art, and to discuss approaches for improvements and innovation. The workshop aims for sharing knowledge about how to apply today's tools most successfully, and for generating a common understanding about what is needed for network research to progress more rapidly and to ensure widely reproducible results. The workshop solicits submissions that improve our understanding of the current state-of-the-art, and that help to identify improved models, methods and tools.



## ACM SIGCOMM 2017 Reproducibility Workshop (Reproducibility'17)

### Call For Papers

Ensuring the reproducibility of results is essential in experimental sciences. Unfortunately, as highlighted recently, a large proportion of research results are hardly, if not at all, reproducible, raising reasonable doubts on the research lead by scientists around the world.

Why do we lack reproducibility?

Why do we lack reproducibility?

Because we lack **incentives**.

# Key players need incentives

**Authors**

**Reviewers**

**Independent  
Reproducers**

# Key players need incentives



**Authors**

Preparation  
takes time

Details may  
conflict with  
reading flow

**Reviewers**

**Independent  
Reproducers**

# Key players need incentives

Authors

Preparation  
takes time

Details may  
conflict with  
reading flow

Reviewers

Are already  
overloaded

Independent  
Reproducers



# Key players need incentives

**Authors**

Preparation  
takes time

Details may  
conflict with  
reading flow

**Reviewers**

Are already  
overloaded

**Independent  
Reproducers**

Reproduction  
papers are  
rarely accepted

# Why does reproducibility **enforcement** not help?

Reject non-reproducible papers.

# Why does reproducibility **enforcement** not help?

Reject non-reproducible papers.



Preparation  
takes time

Details may  
conflict with  
reading flow



Are already  
overloaded



Reproduction  
papers are  
rarely accepted

# Observations

Enforcing reproducibility does not solve all problems.

We need multiple independent reproductions.

Visibility is a major incentive.

We need an **ecosystem** that is based on visibility.

Supportive measures may help.

Supportive measures should adapt to experiences.

# Building blocks for an ecosystem to support reproducible research

# (SIGCOMM) Reproducibility Challenge

## Core idea of an Reproducibility Challenge

- Venue to submit reproduction reports
- Reproducers gain visibility
- Badly reproducible papers would be highlighted, increasing incentives for authors
- Mix of proceedings, presentations, and live “hacking”

## Why at SIGCOMM?

- High visibility, high incentives for both authors and reproducers
- Authors of original papers likely attend SIGCOMM for other reasons

# (SIGCOMM) Reproducibility Challenge

## Core idea of an Reproducibility Challenge

- Venue to submit reproduction reports
- Reproducers gain visibility
- Badly reproducible papers would be highlighted, increasing incentives for authors
- Mix of proceedings, presentations, and live “hacking”

**Most basic building block.**

## Why at SIGCOMM?

- High visibility, high incentives for both authors and reproducers
- Authors of original papers likely attend SIGCOMM for other reasons

# Explicit incentives for authors

## Requirement

- Articulate expectations about reproducibility in Call for Papers

## Possible implementations

- Reproducibility is a tie breaker
- Positive comments at the conference, highlighted in proceedings
- Reproducibility award



# Reproducibility review

## Idea

- First, review technical merit, then, second, review reproducibility

## Problem

- Timeliness

## Approach

- Establish a SIGCOMM Reproducibility Review Committee (RRC)
  - Central pool of proficient graduate students chaired by a senior members
- PCs may submit subset of papers to RRC
  - Papers that are likely to be accepted and claim reproducibility
- RRC may help with reproducibility shepherd
  - Helps authors to select meta data format, storage sites etc.

# Metrics, Badging, Journal Fast Tracking

Requires high experience with reproducibility papers

Building Block	Initial	Evolved	Mature
Reproducibility Challenge	✓	✓	✓
Author Incentives	✗	✓	✓
Reproducibility Review	✗	✓	✓
Metrics & Badging	✗	✗	✓
Journal Fast-Track	✗	✗	✓

# Conclusion

Lack of reproducibility mainly because of lack of incentives

Incentives are needed for authors *and* reproducers

We need a forum for reproducibility

Expose reproducibility and non-reproducibility

Co-locate **Reproducibility Challenge** with **visible** and established conference

# Next steps


Reproducibility Challenge  
@ SIGCOMM 2018?

# Backup

# Comparison with other approaches

ACM Conf. On Hybrid Systems:  
Computation and Control

Artifact Evaluation Committee



Building Block	CCR	AEC	HSCC
Reproducibility Challenge	✗	✗	✗
Author Incentives	✓	n/a <sup>1</sup>	✓
Reproducibility Review	✓	✓	✓
Metrics & Badging	✓	✗	✓
Journal Fast-Track	n/a	n/a <sup>1</sup>	✗

1: Details depend on specific venue, not central AEC.