# *TriScale*: A Framework Supporting Replicable Performance Evaluations in Networking

ANONYMOUS AUTHOR(S)

When designing their performance evaluations, networking researchers often encounter questions such as: How long should a run be? How many runs to perform? How to account for the variability across multiple runs? What statistical methods can be used to analyze the data? Despite their best intentions, researchers often answer these questions differently, thus impairing the replicability of their evaluations and decreasing the confidence in their results.

To support networking researchers, we propose a *systematic methodology* that streamlines the design and analysis of performance evaluations. The hierarchical approach partitions the performance evaluation in a sequence of stages that build on top of each other and follow the principle of separation of concern. We exemplify this basic approach for a large class of performance evaluation scenarios. To this end, we first identify for each stage the temporal characteristics of variability sources. We then apply, for each source, rigorous statistical methods to derive performance results with *quantifiable confidence*, in spite of the inherent variability. We implement an instance of that methodology in a software framework called *TriScale*. For each performance metric, *TriScale* computes a variability score that estimates, with a given confidence, how similar the results would be if the evaluation were repeated; in other words, *TriScale quantifies the replicability* of the performance evaluation. We apply *TriScale* to four different use cases (congestion control, wireless embedded systems, failure detection, video streaming), demonstrating that *TriScale* helps generalize and strengthen previously published results.

Improving the standards of replicability in networking is a crucial and complex challenge; with *TriScale*, we make an important contribution to this endeavor by providing a rationale and statistically sound experimental methodology.

## 1 INTRODUCTION

The ability to replicate an experimental result is essential for making a scientifically sound claim. In networking research, replicability[1] is a well-recognized problem due to the *inherent variability of the experimental conditions*: The uncontrollable dynamics of real networks [21, 51] and the time-varying performance of hardware and software components [15, 49, 72] cause major changes in the experimental conditions, making it difficult to replicate results and quantitatively compare different solutions [8]. In addition, *differences in the methodology* used to design an experiment, process the measurements, and reason about the outcomes impair the ability to replicate results and assess the validity of claims reported by other researchers. Without replicability, any performance evaluation is debatable at best.

To be replicable, performance evaluations must account for the inherent variability of networking experiments on different time scales and therefore, experiments are typically repeated to increase the confidence in the conclusions. To facilitate this, the networking community has put great efforts into developing testbeds [55] and data collection frameworks [81]. However, there is still no *systematic methodology* that specifies how to design and analyze performance evaluations. The literature is currently limited to generic guidelines [9, 52, 63] and recommendations [38, 43, 57], which leave critical questions open *before* an experiment (How many runs? How long should a run be?) and *after* (How to process the data and analyze the results?). Without a systematic methodology, networking researchers often design and analyze similar experiments in different ways, making

[1]Different terminology is used to refer to different aspects of replicability research [12, 59]. In this paper, we refer to replicability as the ability of different researchers to follow the steps described in published work, collect new data using the same tools, and eventually obtain the same results, within the margins of experimental error. This is usually called replicability [1] but sometimes refer to as reproducibility.

them hardly comparable [16]. Yet, strong claims are being made ("our system improves latency by 3×") while confidence is often discussed only in qualitative ways ("with high confidence"), if at all [72, 80]. Furthermore, it is currently unclear how to assess whether a networking experiment is indeed replicable. We argue that a systematic methodology is needed to help to resolve this situation.

To achieve this goal, we face four key challenges that are currently unsolved.

**Rationality** The methodology must rationalize the experiment design by linking the design questions (*e.g.*, How many runs?) with the desired confidence in the results.

**Robustness** The methodology must be robust against the variability of the experimental conditions. The data analysis must use statistics that are compatible with the nature of networking data and be able to quantify the expected performance variation shall the evaluation be replicated.

**Generality** The methodology must be applicable to a wide range of performance metrics, evaluation scenarios (emulator, testbed, real world), and network types (wired, wireless).

**Conciseness** The methodology must describe the experimental design and the data analysis in a concise and unambiguous way to foster replicability while minimizing the use of highly treasured space in scientific papers.

This paper presents the following four contributions.

- We propose a *systematic methodology* that streamlines the design and analysis of performance evaluations and supports the replicability of networking evaluations by addressing the four challenges mentioned above. The method is based on a hierarchical approach that partitions the performance evaluation in a sequence of stages that build on top of each other, following the principle of separation of concern.
- We propose a *concrete instance* of this approach; it is applicable to a large class of performance evaluation scenarios and is based on an analysis of the temporal characteristics of variability sources in networking experiments (§ 2.2). For each variability source, we apply appropriate and rigorous statistical methods to derive performance results with *quantifiable confidence* (§ 3).
- We implement our proposal in an *extensible and publicly available software framework* called *TriScale* [5]. For each performance metric, *TriScale* computes a variability score that estimates, with a given confidence, how similar the results would be if the evaluation were replicated.
- We illustrate the benefits and generality of *TriScale* through *four different case studies* (§ 6) involving both testbed experiments and network emulations: congestion control, wireless embedded systems, failure detection, and video streaming. These case studies demonstrate how the lack of a systematic methodology has led to erroneous or unfair comparisons between protocols or, conversely, that *TriScale* helps generalize and strengthen previously published results and claims.

Most prior work towards replicable networking research focus on data collection [55, 81]. *TriScale* complements these efforts by providing the first general framework that guides networking researchers through the design of their experiments and the analysis of the gathered data, while also *quantifying the replicability* of the performance evaluation.

We strive to make this paper itself "replicable": all data and source code are openly available [5, 7]. Most plots are created using *TriScale* and are interactive: they are hyperlinks to online versions allowing for dynamic visualizations.

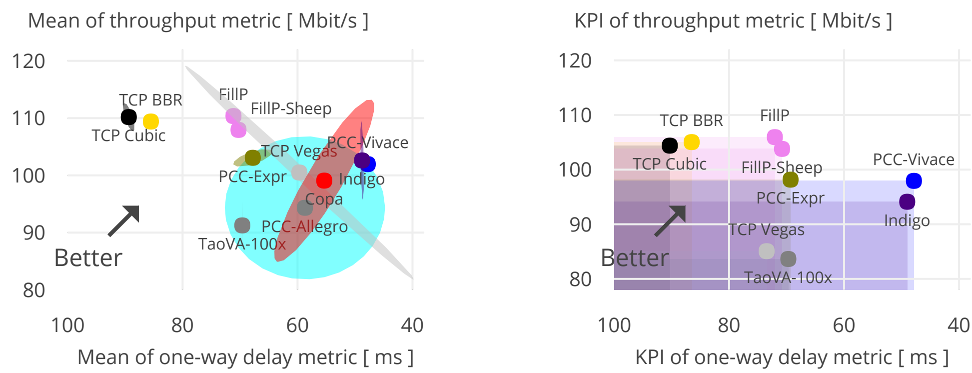


(a) Data analysis with Pantheon [81] (replicated). *Dots represent the mean performance across all runs; metrics are the mean throughput and 95th percentile of the one-way delay; ellipses represent the 1σ performance variation across all runs, where σ denotes the standard deviation.*

(b) Data analysis with *TriScale. Dots represent Key Performance Indicators (KPIs) across all runs: the 25th percentile of the throughput metric and the 75th percentile of the one-way delay metric (same metrics as in Fig. 1a). KPIs are estimated with 75% confidence. Shaded areas represent dominance regions.*

Fig. 1. Sample data from the congestion-control case study (§ 6.1). *The same data may be analyzed in different ways. Compared with Pantheon's analysis (Fig. 1a), TriScale's analysis allows for a more intuitive interpretation of the results (Fig. 1b): The performance of each scheme is reduced to a single point, TriScale's KPIs, which makes the comparison between the schemes unambiguous. These KPIs are not arbitrary: they are robust non-parametric statistics estimating, with a given confidence level, the expected performance if the experiment was repeated. Thus, TriScale's KPIs inherently account for the variability in the results.*

## 2 OVERVIEW OF *TRISCALE*

This section illustrates how *TriScale* improves the analysis of experimental results with a concrete example (§ 2.1) then presents the core principles of the methodology (§ 2.2).

### 2.1 How *TriScale* Improves Data Analysis

Assume you are new to the field of congestion control and would like to understand the strengths and weaknesses of the state-of-the-art. Luckily, the community develops useful tools like Pantheon [81], a data collection framework that facilitates comparisons of congestion-control schemes.

You are particularly interested in throughput and one-way delay of full-throttle flows, *i.e.*, flows whose performance is only limited by the congestion control. You start with one flow and evaluate performance using MahiMahi [53], a traffic and network emulator integrated in Pantheon, using the same settings as in [81]: 10 runs of 30 seconds each. You collect data for the 17 congestion-control schemes available.

Pantheon assists you in collecting the data, but not in their analysis or interpretation. Yet, these are two non-trivial tasks. Consider, for example, the results shown in Fig. 1a (replicated from [81]); the dots represent mean performance across all runs using mean throughput and 95th percentile of the one-way delay as metrics. Multiple questions arise:

(Q1) Can the schemes be compared? It appears that *TCP Vegas* performs better than, *e.g.*, *TaoVA-100x*. However, the ellipses capture the variability of the results across all runs; more precisely, they represent the $1\sigma$ variation across runs, where $\sigma$ is the standard deviation. What can you then conclude about the actual performance of these schemes? Can you conclude anything if the ellipses are overlapping? *E.g.*, can you say that *TCP Vegas* performs better than *PCC-Expr*?

(Q2) What is the confidence in the comparison? Intuitively, the results of, *e.g.*, *PCC-Allegro*, which have a large variability, are less trustworthy than those of, *e.g.*, *FillP-Sheep*, for which the ellipse is hardly visible. But can you quantify the confidence in this result?

(Q3) Is a runtime of 30 seconds sufficiently long to fairly compare the different schemes?

These questions relate to the robustness and rationality challenges (§ 1) and are left unanswered by the analysis shown in Fig. 1a. In fact, the analysis may suggest wrong interpretations. Ellipses are a two-dimensional representation of the standard deviation across runs, suggesting that one can expect about 68% of the data points to fall in that region. However, this is correct *only if* the underlying distribution is normal, which is hardly ever true (§ 3).

Fig. 1b illustrates the *same data* analyzed with *TriScale*. The dots now represent *TriScale*'s Key Performance Indicators (KPIs). A KPI estimates a given percentile of a performance metric's underlying distribution (*i.e.*, the unknown distribution we would obtain with infinitely many samples) with a certain confidence. We use the same performance metrics: the mean throughput and 95th percentile of the one-way delay, for which we have 10 samples (one per run). Based on these 10 samples, *TriScale* estimates the 25th percentile of the throughput metric (*higher* throughput is better) and the 75th percentile of the one-way delay metric (*lower* delay is better) We choose a 75% confidence level for the estimation of both KPIs.[2] In other words, with a 75% confidence, 75% of the runs yield a performance that is as least as good as the KPI values (*i.e.*, equal or higher throughput and equal or lower one-way delay). Note that, in this paper, we consider multiple performance dimensions (*e.g.*, throughput and delay) independently. The approach can be extended towards multi-objective performance evaluations using the principles of Pareto-dominance, but such extension is beyond the scope of this paper.

Using this methodology (detailed in § 2.2 and § 4), *TriScale* answers the three open questions mentioned previously:

(Q1) Since the KPIs are individual dots, we can unambiguously compare different schemes with respect to the two performance metrics. Contrary to what Fig. 1a suggests, we observe in Fig. 1b that *TCP Vegas* is not strictly better than *TaoVA-100x*, as *TCP Vegas* performs worse terms of one-way delay; also, *PCC-Expr* performs better than *TCP Vegas* in both performance metrics.

(Q2) The KPIs confidence levels explicitly state how confident we are with these results.

(Q3) *TriScale* tests whether the different schemes have converged (§ 4.5), *i.e.*, have the metrics reached stable values within the experiment runtime?

**Summary.** Tools like Pantheon [81] support data collection, but leave the design of the experiments and the data analysis up to the researcher, leading to ambiguous interpretations and non-replicable results. *TriScale* aims to fill this gap.

## 2.2 Core Principles of *TriScale*

*TriScale* is a framework for networking experiments (Fig. 2); it is based on a systematic methodology that streamlines the design and analysis of performance evaluations to improve the replicability of networking evaluations. Its hierarchical approach partitions the performance evaluation in a

[2] In this example, we use a confidence level of 75% which is a low value (95% would be more common). However, estimating 25th and 75th percentiles with 95% confidence requires at least 11 data points (Eq. (3)), while Pantheon currently performs series of 10 runs. To compare *TriScale*'s and Pantheon's analysis methods, we chose here to lower the confidence level and keep the same number of samples.

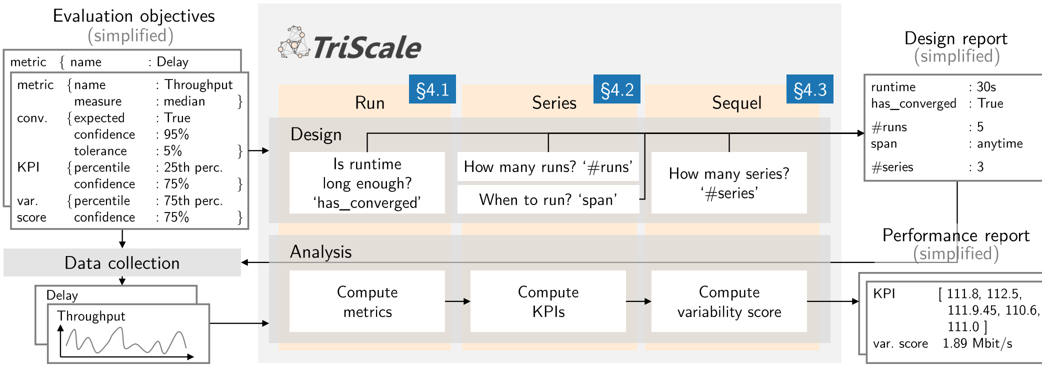


Fig. 2. Overview of *TriScale. TriScale is a framework supporting the design and analysis of networking experiments. TriScale assists the user in the design phase with a systematic methodology to answer important experiment design questions such as "How many runs?" and "How long should the runs be?". After the raw data are collected, TriScale supports the user by automating the data analysis. The framework implements robust statistics that handle the intrinsic variability of experimental networking data and returns expressive performance reports along with a variability score that quantifies the replicability of an experiment.*

sequence of stages that build on top of each other and follow the principle of separation of concern. Specifically, it separates the entire evaluation into three timescales, hence the name of ***Tri**Scale*.

Given the user's objectives (e.g., the KPIs to analyze and the confidence levels to reach), *TriScale* helps to answer questions such as: How many runs should be done? How long should the runs be? When to perform the runs? Based on these answers, the user can then proceed with the data collection. In the analysis phase, the user provides those data to *TriScale* which automatically produces expressive and easy-to-interpret performance reports together with variability scores that quantify the replicability of the experiment.

It is important to note that the specific models and methods within each of the timescales are dependent on the specific class of systems and performance evaluations that will be undertaken. In this paper, we provide a concrete instance of *TriScale* that is applicable to a large class of networking performance evaluations, exemplified by the four different case studies (§ 6): congestion control, wireless embedded systems, failure detection, and video streaming.

In the following, we explain *TriScale*'s main building blocks. We start by describing the three timescales underlying the methodology, then describe how *TriScale* concretely supports the users with their experiments design and analysis.

**Timescales.** We structure *TriScale*'s methodology around three timescales: runs, series (of runs), and sequels (of series). These timescales intuitively capture the different sources of variability underlying performance evaluations in networking.

A *run* is one execution of an evaluation scenario, *e.g.*, a 30 s execution of *TCP BBR*. During a run, some performance dimensions are measured, *e.g.*, packet delay, which vary due to different sources of variability such as protocol dynamics and cross-traffic. The performance during a run is summarized by a *metric*, for example the 95th percentile. Depending on the scenario, one may want the metric to estimate long-term performance, for example in case of long-lasting flows. The run should then be sufficiently long to let the metric value converge.

Typically, one executes multiple runs to measure performance under different conditions. We call such a set of runs a *series*. For example, we may execute 100 runs within one week, from which we obtain a set of metric values, one for each run. We summarize the performance throughout a series with a *key performance indicator (KPI)* that estimates the expected performance for any possible run in the time span of the series, for example by estimating the median of the metric distribution. The intuition is that, with a series of run, one randomly samples the distribution of possible experimental conditions during that week, which allows to estimate the distribution of the performance metric of interest.

In general, variability sources such as cross-traffic vary with an a priori unknown temporal long-term correlation; in other words, the distribution of conditions sampled during a series may not be stationary but time-varying. Therefore, in order to generalize the results, one should perform multiple series, which we call *sequels*. Intuitively, sequels allow to estimate the expected performance for any series, for example the expected KPI for any week. Our method uses sequels to compute a *variability score* that serves to *quantify the replicability* of an experiment. To this end, it defines a confidence interval for the expected results one would obtain shall new series of runs be performed.

*TriScale* uses these three timescales of runs, series, and sequels to divide the experiment design and data analysis pipelines.

**Experiment design.** The design phase starts with the definition of the evaluation objectives (Fig. 2, left). For each performance dimension, the user defines the metric, the convergence requirements, a KPI, and a variability score (§ 4). Please note again that this paper simply deals with multiple performance dimensions independently of one another. Given these inputs, *TriScale* derives the minimum number of runs (*#runs*) and series (*#series*) needed to compute the chosen KPIs and variability scores, thus answering the question of how many runs to perform. Using data from test runs or previous experiments, *TriScale* can assess whether the runtime appears long enough to let the metric value converge. Additionally, *TriScale* can make use of these test runs to identify time-dependent patterns in the experimental conditions (§ 4.6). This is important to understand the root cause of the statistical behavior of the measurements, and helps to answer the question of *when* the runs should be performed. Note that the congestion-control example presented previously uses network emulation; thus, there is no time dependency and it does not matter when the experiment is performed (*i.e.*, *span: anytime*). The design phase results in a report (Fig. 2, right) summarizing how to run the experiments. Based on this report, the user can collect the raw data then moves on to the analysis phase.

**Data analysis.** Once the experiment has been designed and the data collected, the raw data are passed to *TriScale* for a three-stage analysis, one per timescale. First, the raw data from one run are processed, *i.e.*, convergence is assessed and the performance metrics are computed. Then, the performed series allow accounting for the short-term variability in the experiments. This timescale leads to one number per series and per metric: the KPIs (§ 4.2). Finally, the sequels (repetition of series) are used to compute a *variability score* capturing the long-term variability of the KPIs. This timescale leads to one number per metric (§ 4.3).

**Using *TriScale*.** *TriScale* is implemented as a Python module [5], as detailed in § 5. For each timescale, a dedicated function performs the corresponding test or analysis. The functions take as input raw data in the form of CSV files, Pandas DataFrames, or Python arrays; the outputs are returned and optionally saved as CSV files. These same functions also produce data visualizations such as those shown in Fig. 3 to 5.

We try to make *TriScale* intuitive and easy to use. For a better impression of *TriScale*'s usability, an interactive demo is available and can be run directly in your web browser [6].

## 3 STATISTICS FOR REPLICABILITY

This section briefly reviews classes of statistical approaches and motivates the choice of the methods we use in the current implementation of *TriScale* to handle the variability inherent to networking evaluations. Please note that we take a safe evaluation approach in the described instance of the *TriScale* approach: we do not suppose any knowledge about the statistical distributions underlying the variability of the measurements. Of course, tighter estimates are possible if additional reliable information is available.

**Descriptive and predictive statistics.** A statistic is a number computed from some data using a mathematical formula; it can always be calculated and provides a factual description of the underlying data. This is referred to as a *descriptive statistic*. In addition, certain statistics have some *inference* power; *i.e.*, based on the collected data, one may infer the shape of the (unknown) underlying data distribution. These are then referred to as *predictive statistics*.

Predictions are always uncertain and rely on certain hypotheses. If the hypotheses hold for the collected data, then predictive statistics estimate, with a quantifiable level of confidence, some property of the underlying distribution such as the mean or the median. One can then predict expected values of data samples that have not been collected. A common hypothesis is that the collected data is *independent and identically distributed* (*i.i.d.*). Informally, this means that the underlying distribution of the data does not change and that successive data samples are uncorrelated. It is also common to presume the *nature of the data distribution*, for example normal or Poisson distribution. For example, one can estimate the mean $\mu$ and standard deviation $\sigma$ of a distribution based on a data sample. If the underlying data distribution is normal (the hypothesis), we can infer that about 68% of all data points will be contained within $\mu \pm \sigma$ (the prediction). But if the distribution is *not normal*, the statistics $\mu$ and $\sigma$ are *only* descriptive, *i.e.*, they do not predict anything about unseen samples.

**Statistical methods.** There are two common classes of statistical approaches: hypothesis testing and estimation. *Hypothesis testing* consists of formulating a so-called null hypothesis that the test aims to reject. Based on the collected data, one computes the probability, called the $p$-value, that the null hypothesis is correct. If the $p$-value is sufficiently low, the null hypothesis is rejected and considered proven incorrect. For example, the one-way ANOVA [77] is a common method to test for significant differences in the mean of multiple data samples. *Estimation* consists of computing confidence intervals (CIs) for a given parameter (*e.g.*, the mean of a distribution). A CI is always associated with a certain confidence level (*e.g.*, a 95% CI) which is the probability that the interval includes the true value of the parameter. For example, $[a, b]$ is a 95% CI for the mean if the true mean value is between $a$ and $b$ with a probability of at least 95%.

These approaches are further classified as *parametric* if the nature of the underlying distribution is known and as *non-parametric* if no assumptions are made about the underlying distribution. For example, the Kruskal-Wallis test [76] is the non-parametric equivalent of the one-way ANOVA. The tests are similar, but the former does not assume that the underlying distribution is normal. The central limit theorem [79] offers another alternative to handle unknown distributions, but it only allows to argue about the arithmetic mean.

**Statistics for replicability in networking.** Informally, replicability is the principle that the "same experiment" leads to the "same results." Assessing replicability entails predicting whether future data (the results of a newly-performed experiment) will be the same as the known data (the results of previous experiments); it is a prediction. *TriScale* formalizes this intuitive notion by introducing the variability score, see § 4.3, § 4.4, and § 4.7.

Literature reports that experimental data are rarely normal [49, 64] and hence recommends using *non-parametric* statistics. One should also consider *robust statistics*, for example using median instead of mean, that are not overly skewed by outliers, as these are common in networking data. While hypothesis testing is commonly used, statisticians argue that the methods are misunderstood and misused [44] and are thus calling for a change in scientific practices [27, 74]. We favor *estimation* over hypothesis testing because CIs are more legible than *p*-values and easier to interpret. Furthermore, the confidence level of an estimation only depends on the sample size, which is useful to guide the experimental design.

In 1936, Thompson introduced a method to compute non-parametric CIs for percentiles [71]. This approach is found in statistics [28] and computer science [45] textbooks, but it is barely used today ([13, 49, 64] are the few exceptions). As Thompson's method is well-suited to handle the variability of experimental networking data, we use it in the described instance of *TriScale*'s methodology (§ 4.5). In this paper, we illustrate the potential of the approach (§ 6) and strive to facilitate its use by providing the necessary software support (§ 5).

## 4 DESIGNING *TRISCALE*

In this section, we first describe the data analysis performed by *TriScale* and how the analysis procedure is linked to the design of experiments (§ 4.1 to § 4.3). We then illustrate how the formalism introduced by *TriScale* allows to unambiguously describe an entire performance evaluation with only a handful of parameters (§ 4.4). Thereafter, we detail the robust and non-parametric statistical methods used by this instance of *TriScale* (§ 4.5), and discuss how the framework assists a user in deciding the required time span for a series of runs (§ 4.6). We finally show how *TriScale* helps in assessing the replicability of experiments by computing a variability score (§ 4.7).

### 4.1 Runs and Metrics

In *TriScale*, metrics evaluate a performance dimension across a run, for example, the mean throughput achieved by a congestion-control scheme over 30 s runtime of a full-throttle flow. Computing a metric takes the following inputs.

**Inputs.**
- The metric *measure*, *e.g.*, mean, maximum;
- The *convergence* requirements
  { expected : true/false ,
  confidence : $C$ (*default: 95%*) ,
  tolerance : $t$ (*default: 5%*) };
- The raw data of the run.

In general, any measure can be used. The current implementation of *TriScale* (§ 5) supports the arithmetic mean, the minimum, the maximum, and any percentile.

**Procedure.** If the run is expected to converge, *TriScale* starts by performing a convergence test, whose purpose is to assess whether the metric has reached a stable value by the end of the run (and if it is therefore a reliable estimate of the long-running performance). Note that the performance dimensions and convergence behavior can largely vary between the systems. Therefore, suitable methods to test for convergence may vary and need to be considered during the design of an experiment. In the current instance of *TriScale*, we implement an approach that is well-suited to our networking case studies (§ 6).

The implemented convergence test procedure starts by computing metric values over a sliding window of the raw data points, with a fixed size of half the data points. For each window, one metric value is computed, starting with the first half of the data. The window repeatedly slides by 100th of the number of samples until all data points are used, leading to a set of 100 metric values.

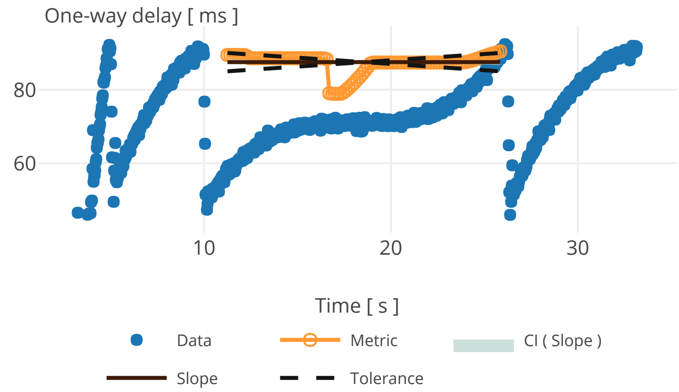


(a) Raw data (one-way delay) and metric (95th percentile). *Example run of TCP Cubic.*

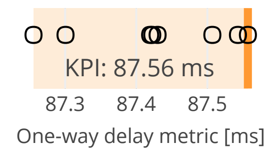


(b) Runs' metric data and corresponding KPI value.

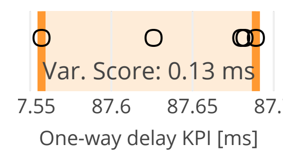


(c) Series' KPI data and corresponding variability score.

Fig. 3. Example plots produced by *TriScale* during the data analysis. *Fig. 3a: computation of the metric (95th percentile on one-way delay) with convergence test (confidence 95%, tolerance 5%). Fig. 3b: computation of the KPI (75th percentile with 75% confidence). Fig. 3c: computation of the variability score (25-75th percentile range with 75% confidence). Sample data from the case study in § 6 for TCP Cubic. Plots are produced by TriScale.*

*TriScale* performs its convergence test (detailed in § 4.5) on the metric values. This procedure tests the convergence of the *metric*, and not of the raw data, in order to reduce the impact of transient behavior on the convergence test. If the test is passed, *TriScale* returns the median of the converged metric values as run metric. If convergence is not expected, *TriScale* simply computes the run metric over the whole raw data.

**Outputs.**
- The result of the convergence test (if any);
- The metric value for the run;
- Textual logs, plot of the input data, and metric.

**Link to the experiment design.** The computation of metrics is linked to the definition of the *runtime*, *i.e.*, how long a run should be and which of its raw data to take. If the evaluation scenario is terminating, for example transmit 1 MB of data, the runtime must be long enough to complete the task. If the evaluation is long-running, for example the estimation of operational lifetime, the runtime must be long enough for the metric (*e.g.*, energy consumption) to converge. Details about the specific convergence test are described in § 4.5. *TriScale* can analyze experiments to estimate whether the runtime appears long enough as illustrated in § 6, *i.e.*, it can assess with quantifiable confidence that the metric values are stable for a certain runtime, under certain assumptions on the global convergence behavior. However, *TriScale cannot guarantee* that the runtime is long enough for a sound evaluation of long-running performance, as this requires context-specific knowledge.

### 4.2 Series and KPIs

*TriScale*'s key performance indicators (KPIs) evaluate performance dimensions across a series of runs. Performing multiple runs allows to mitigate the inherent variability of the experimental conditions. KPIs capture this variability by estimating percentiles of the unknown metric distributions. Concretely, a *TriScale* KPI is a one-sided CI of a percentile, *e.g.*, a lower bound for the 75th percentile of the throughput metric, estimated with a 95% confidence level.

**Inputs.**
- The KPI definition
  { percentile : $p$ ,
  confidence : $C$ };
- The metric data from a series of runs.

**Procedure.** To compute KPIs (*i.e.*, to compute a CI for a given percentile), *TriScale* uses Thompson's method (§ 4.5) which requires the input data to be *i.i.d.*. *TriScale* starts by performing an independence test (§ 4.5) to check that the metric data empirically appears *i.i.d.* before computing the KPI.

**Outputs.**
- The result of the independence test;
- The KPI value for the series of runs;
- Textual logs, plot of the metric data and its KPI.

**Link to the experiment design.** The computation of KPIs is linked to the definition of the number of runs in a series (*# runs*) and the series time span (*span*). The minimal number of runs in a series directly follows from the definition of the KPI, *i.e.*, the percentile to estimate $p$ and the desired confidence level $C$ (see Eq. (3)). The series time span refers to the time interval used for scheduling the runs in a series, *i.e.*, when to run the experiment. This is important because networks often feature time-dependent conditions; for example, there may be systematically more cross-traffic during daytime than nighttime. Failing to consider such dependencies may bias the results and yield wrong conclusions. This concept of series also applies when "slicing" a long experiment into smaller independent ones. In such a case, it is crucial to consider warm-up and cool-down effects to avoid biasing the results. Note that such slicing strategy is more likely to result in empirically non-*i.i.d.* data than a random schedule of truly independent runs. *TriScale* helps to detect certain classes of dependencies with a dedicated "network profiling" function (example in § 6). Here again, other dependency analysis methods can simply be added in order to better tailor *TriScale* to a specific class of systems under evaluation.

### 4.3 Sequels and Variability Score

*Sequels* are repetitions of series of runs. *TriScale*'s variability score evaluates the variations of KPI values across sequels. Sequels enable *TriScale* to detect long-term variations of KPIs and ultimately quantify the replicability of an experiment.

Concretely, a variability score is made of two one-sided CI for a symmetric pair of percentiles. For example, a 75% confidence interval for the 25-75th percentile range of the delay KPIs from all sequels. Again, we attach a confidence value to the confidence interval, or equivalently, to the percentile estimations.

**Inputs.**
- The variability score definition
  { percentile : $p$ (or 1-$p$),
  confidence : $C$};
- The KPI values of each sequel.

**Procedure.** The procedure is the same as for the KPIs: Since Thompson's method requires the input data to be *i.i.d.* (§ 4.5), *TriScale* first performs an independence test on the KPI data before computing the variability score.

**Outputs.**
- The result of the independence test;
- The variability score value for the entire sequels;
- Textual logs, plot of KPI data, and corresponding variability score.

Table 1. Exemplary evaluation parameters of typical for networking (med=median). * *TriScale returns the minimal number of runs (#runs) and series (#series) based on the definition of KPI and variability score, respectively.*

| **Use case** | **Evaluation Objectives** | | | | | | | | **Experiment Design** | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Metric | Convergence | | | KPI | | Var.Score | | | | | |
| | Measure | Exp. | Conf. | Tol. | Perc. | Conf. | Perc. | Conf. | #runs* | #series* | runtime | span |
| Latency of | | | | | 95 | 95% | med | 75% | 59 | 3 | | |
| real-time | max | True | 95% | 5% | 95 | 99% | 75 | 75% | 90 | 5 | Depend on | |
| protocol | | | | | 99 | 95% | med | 90% | 299 | 5 | networks and | |
| Average | med | False | - | - | 90 | 95% | med | 90% | 29 | 5 | protocols | |
| delay | $90^{th}$ perc. | | | | med | 95% | med | 90% | 5 | 5 | | |

**Link to the experiment design.** The computation of the variability score is linked to the definition of the number of series (*#series*). The minimal number of series directly follows from the definition of the variability score, *i.e.*, the percentile to estimate $p$ and the desired confidence level $C$ (Eq. (3)).

## 4.4 Formalism Brings Conciseness

*TriScale* formalizes the definition of the evaluation objectives. As illustrated in Fig. 2, for each performance dimension, the user defines a metric with convergence requirements, a KPI, and a variability score. *TriScale* links these objectives with the experiment design, resulting in four additional parameters: the number of runs per series (*# runs*), the number of series (*# series*), the length of a run (*runtime*), and the time span of a series (*span*).

Thanks to this formalism, *TriScale* solves the *conciseness* challenge: altogether, these 12 parameters are sufficient to formally describe the entire performance evaluation. Since the data analysis in *TriScale* is automated and deterministic, documenting these parameters guarantees computational reproducibility, *i.e.*, the ability to recreate the results when all raw data are available [47].

Table 1 shows a few examples of concrete parameter settings for typical networking evaluation use cases. For example, evaluating the latency of a real-time protocol requires high confidence levels for extreme percentiles. This quickly increases the number of runs that must be performed, *e.g.*, at least 90 for estimating the 95th percentile with 99% confidence and at least 299 for estimating the 99th percentile with 95% confidence. This illustrates that it is "easier" to increase the confidence level of an estimation than to estimate a more extreme percentile with the same confidence level. Note that both *#runs* and *#series* are only derived from the definition of the KPI and variability score; *i.e.*, these parameters are not influenced by the runtime or the time span of an experiment.

The second use case in Table 1 (bottom rows) illustrates two different perspectives on "averages" using delay as an example. If the metric is the median and the KPI is the 90th percentile, one can conclude that 90% of the runs have a median delay equal or better than the KPI value. Conversely, if the metric is the 90th percentile and the KPI is the median, one can conclude that, in half of the runs, the 90th percentile of the delay in the run is equal or better than the KPI. Both are "averages", but with different meanings and different requirements in terms of number of runs. Only users can know what is more appropriate for their own use case.

## 4.5 Statistics in *TriScale*

As discussed in § 3, performance evaluations in *TriScale* focus on statistics that are both robust (*i.e.*, tolerant to outliers) and non-parametric (*i.e.*, making no assumption about the nature of the data distribution). Of course, one can implement in *TriScale* methods resulting in tighter estimates if additional reliable information about the underlying distribution of the data is available.

In its current implementation, *TriScale* uses three carefully-chosen statistical methods. We first present the convergence test used in the computation of metrics, which is based on the Theil-Sen linear regression [68, 70]. We then introduce the computation of confidence intervals using Thompson's method [71]. Since this method requires the data to be *i.i.d.*, *TriScale* checks empirically whether this requirement is satisfied with an independence test, presented last. Finally, we discuss the consequences of failing tests for the collected data.

**Convergence test.** When an evaluation aims to estimate the long-running performance, *i.e.*, the expected performance if the run would continue for a very long time, one must verify whether the runs are long enough to produce reliable estimates. To verify this, *TriScale* implements a convergence test based on the Theil-Sen linear regression [68, 70]. This approach computes the slope of the regression line as the median of all slopes between paired values. A $C\%$ confidence interval (CI) for the slope is defined as the interval containing the middle $C\%$ of slopes between single pairs. *TriScale*'s convergence test is passed if the $C\%$ CI for the regression is included in the tolerance value ($\pm\, t\%$). To test the convergence of a run, *TriScale* uses the confidence $C$ and the tolerance $t$ specified in the evaluation objectives; by default, $C$ and $t$ are set to 95% and 5%, respectively.

Such a test is sensitive to the scale of the input data. To remove this dependency, *TriScale* first maps the data to $[-1, 1]$ using a linear transformation, then performs the convergence test on the scaled data. Hence, the convergence test becomes dimensionless and the same tolerance value can be used to compare different evaluations without bias. Fig. 3a shows an example of the Theil-Sen slope (brown, solid), its CI (light blue, solid), and the tolerance (black, dashed).

Note that this convergence test is based on some assumptions; *e.g.*, that the convergence of metric values is captured by the convergence of the slopes toward zero. This would not hold if one would measure *e.g.*, energy consumption since it is cumulative other time; power draw should be used instead. The same holds true for the normalization using a linear transformation; a logarithmic scale may be more relevant in certain scenarios. Alternative convergence tests can be implemented for these cases.

**Confidence intervals.** *TriScale* defines KPIs and variability scores based on CIs for distribution percentiles, which can be computed using a robust and non-parametric approach based on Thompson's method [71], which has been later on shown to be valid for any independent samples of a continuous distribution [28].

Let us denote by $P_p$ the $p$-th percentile of a distribution and by $\mathbb{P}(X)$ the probability of an event $X$. By definition, every data sample $x$ is smaller than $P_p$ with probability $p$ (and larger with probability $1-p$). For a sorted list of *i.i.d.* samples $x_i$ (where $i = 1..N$), the probability that $P_p$ lies between two consecutive samples follows the binomial distribution [71]:

$$\mathbb{P}(x_k \le P_p \le x_{k+1}) = \binom{N}{k} p^k (1-p)^{N-k}, \quad k = 0..N \tag{1}$$

where we assume $x_0 \to -\infty$ and $x_{N+1} \to +\infty$. From this result it follows that the probability of $P_p$ being larger than any sample $x_m$ ($1 \le m \le N$) can be computed as:

$$\mathbb{P}(x_m \le P_p) = 1 - \sum_{k=0}^{m-1} \binom{N}{k} p^k (1-p)^{N-k} \tag{2}$$

These probabilities are symmetric, that is, $\mathbb{P}(x_m \le P_p) = \mathbb{P}(x_{N-m+1} \ge P_{1-p})$. Eq. (2) provides the upper and lower bounds required for computing the CIs. Further, one can derive the minimum number of samples $N$ needed to compute a CI for any percentile $p$ with any confidence level $C$ [64]:

$$Eq. (2) \quad \Rightarrow \quad N \ \geq \ \frac{\log(1-C)}{\log(1-p)} \tag{3}$$

which defines the minimum number of runs and series required based on the definitions of the KPIs and the variability scores. If the probability distribution is discrete, Eq. (2) becomes an inequality ($\mathbb{P}(x_m \leq P_p) \geq \ldots$ ) that provides a safe (*i.e.*, conservative) estimate of which sample $x_m$ is the bound of the CI of interest [28].

This approach provides robust estimates for distribution percentiles and *does not make any assumption on the nature of the underlying data distribution.* It does, however, require that the data samples are *i.i.d.*: *TriScale* checks whether this requirement holds with an independence test, described next.

**Independence test.** Estimating the percentile of a distribution requires often, if not always, that the samples are *i.i.d.*. This is also the case for Thompson's method [71]. *TriScale* implements an empirical independence test to check whether we can safely treat the samples as *i.i.d.*[3] This independence test is applied to the metric data (resp. KPI data) before the computation of a KPI (resp. a variability score). This poses the particular challenge that the number of data samples may be very small (*e.g.*, 3 or 5 KPI values). *TriScale*'s independence test must therefore not be too strict.

The test proceeds in two steps. First, *TriScale* tests whether the data are *weakly stationary* (*i.e.*, no trend and constant autocorrelation structure [20]). *TriScale* verifies this empirically using its convergence test with a confidence of 50% and a tolerance of 10%; these "loose" parameters are used to compensate for (very) small sample sizes. Second, *TriScale* computes the *sample autocorrelation coefficients*, denoted by $\widehat{\rho_k}$, which measure the linear dependence between values of a weakly stationary data series, where $k$ is the lag between data points. A series of size $N$ is *i.i.d.* with 95% probability if $|\widehat{\rho_k}| \leq 1.95/\sqrt{N}$ for $k \geq 1$ [20].

**What if the tests fail?** The user is responsible for designing the evaluation in such a way that the collected data will (likely) pass the tests. *TriScale* facilitates this by guiding the choice of runtime to pass the convergence test and informing about any network time dependencies (§ 4.6) to pass the independence test. Yet, the data may still be correlated or unstable, leading to failing tests (see examples in § 6). Even in such cases, the data still contain useful information. *TriScale*'s metrics, KPIs, and variability scores can be computed; however, since the corresponding hypotheses do not hold, the statistics are *only descriptive* (§ 3). In other words, they do not predict the expected performance and in particular they cannot (and should not) be used to assess the replicability of the evaluation.

### 4.6 Network Profiling

*TriScale* assists the user in deciding on the time span for a series of runs, *i.e.*, the time interval containing all the runs of one series. This is important in order to avoid biasing the evaluation results with time dependencies in the experimental conditions. Indeed, it is common for real-world networks to exhibit periodic patterns. For example, there may be a lot more cross-traffic (*i.e.*, interference) at specific times of the day. In the statistics literature, these patterns are called *seasonal components*. Neglecting these may bias experiments and lead to wrong conclusions, as we illustrate in § 6 and shown *e.g.*, in [72].

[3] Generally, independence results from the experiment design. For networking experiments, however, it is generally not possible to guarantee independence, *e.g.*, the experimental conditions cannot be fully controlled and may be correlated. In such cases, it is common to empirically check whether the data are correlated. If the effective dependence between data samples is sufficiently low, it is considered safe to treat the samples as *i.i.d.*.

*TriScale*'s network profiling functionality allows to compute the autocorrelation coefficients of link quality data (*e.g.*, [39]). Peaks in the autocorrelation plot suggest seasonal components in the network conditions (see Fig. 5), which helps detect (sometimes unexpected) time dependencies. To avoid biasing the results, the span of a series of runs should be chosen as a multiple of the seasonal components. The same care must be taken when choosing the times when to execute a run within a series; the advisable strategy is to randomly sample the entire span of a series.

### 4.7 Assessing Replicability

Replicability refers to the ability of obtaining "the same" results when performing "the same" experiment. In statistics, such property can be investigated using *equivalence testing* [44], which checks whether the values of some parameter of interest, for example the median, obtained for different samples are sufficiently close to be considered "the same." Unfortunately, there is no general way to define "the same" or even "sufficiently close." One must define in advance a threshold for the equivalence test based on expertise. Then, how to assess replicability of networking experiments? How to design a "replicability test" that fairly adapts to different networking contexts and metrics? After several failed attempts, we conclude that defining a generic threshold for equivalence testing in networking might not be possible. But it may also be not necessary.

We argue that the most important is to confidently estimate the variability of the results, which *TriScale* computes with its variability score (§ 4.3). This score *quantifies replicability*: the larger the score, the less replicable are the results (see the example in § B.1). Shall a binary cut between "replicable" and "not replicable" be desired, a threshold can be set based on the variability score, *e.g.*, "Results are said replicable when the variability score is less than 20 Mbps". Such a threshold can only be context-specific. Thus, deciding on threshold values is more related to benchmarking and therefore falls beyond the scope of *TriScale*, as we discuss in § 7.

## 5 IMPLEMENTATION AND SCALABILITY

### 5.1 All-Included Software Package

One obstacle to the adoption of non-parametric statistics is the lack of support in current scientific libraries; in particular, we had to implement the computation of CIs using Thompson's method ourselves. We implemented *TriScale* as a Python module including all necessary functions to apply our methodology. *TriScale*'s API contains one function for each timescale of the data analysis, with docstrings containing detailed information about each function's usage. The module also includes support tools, such as functions producing visualizations. *TriScale* uses Plotly [60] to create interactive plots in which one can zoom in and out, toggle the visibility of individual traces, read data values on hover, etc. Most plots in this paper have been produced using *TriScale* and all are "clickable": figures are hyperlinks leading to dynamic versions of the plots. Our implementation is open source [5].[4] We use Binder [40] to provide an interactive demo of *TriScale* that runs directly in your web browser [6].

### 5.2 Scalability of *TriScale* Data Analysis

We evaluate the scalability of *TriScale* by measuring its computation time, *i.e.*, we evaluate how the time needed for the data analysis scales with a growing input size. To this end, we only consider the time required for performing computations, and exclude other outputs such as logs and plots (*e.g.*, Fig. 3a). Details are presented in Appendix § A.

[4]The repository is currently anonymous for double-blind review. Everything will be packaged and published on PyPI shall the paper be accepted.

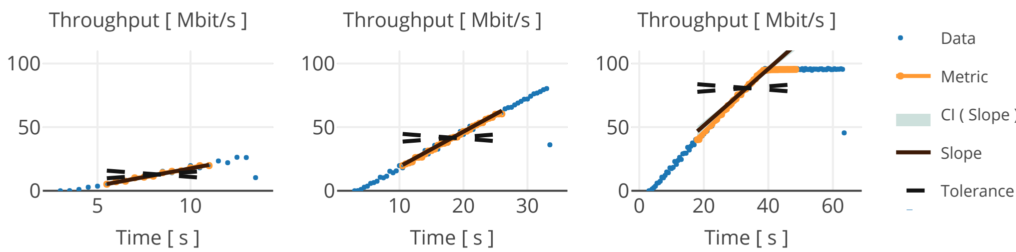


(a) 10 s runtime

(b) 30 s runtime

(c) 60 s runtime

Fig. 4. Egress throughput of the *LEDBAT* congestion-control scheme in MahiMahi [53]. *A runtime of 30 s is clearly not sufficient for LEDBAT's throughput to converge (Fig. 4b). The scheme does converge eventually (Fig. 4c), but even with 60 s runtime, TriScale's convergence test fails: the impact of the start-up phase is too important. See Appendix § B.1 for further details.*

**Conclusion.** The computation time for the data analysis in *TriScale* scales linearly with the input size (Table 2). It is fast (less than 1 s for one million data points on a commodity laptop) and negligible compared to the data collection time.

## 6 *TRISCALE* IN ACTION

We now present four case studies illustrating shortcomings in performance evaluations that *TriScale* addresses, and show how *TriScale* allows generalizing performance claims with a quantifiable confidence. Further details on these case studies (*e.g.*, data, additional plots) are available in § B.

### 6.1 Congestion Control

The first case study shows that, for estimating long-running performance, it is important to carefully set the length of runs (the runtime) and to check whether the performance has converged for the system under evaluation.

We continue the evaluation introduced in § 2.1, which compares congestion-control schemes using Pantheon [81]. Assume we are now interested in *long-running flows*; that is, our goal is to estimate the performance one would obtain if the flows ran "forever". *TriScale*'s convergence test (§ 4.1) checks whether the length of a run is long-enough to provide a robust estimate. Since all schemes are different, it is hard to know a priori the minimum runtime for which the schemes actually converge. For this reason, we test runtimes from 10 to 60 s and check when the schemes pass the test.

For a runtime of 30 s (the one used by the maintainers of Pantheon [56]), only 11 out of 17 schemes pass the test (*i.e.*, converge in most of the cases). *Verus*, *PCC-Allegro*, and *Copa* only converge in less than half of the runs (see § B.1), whereas *QUIC Cubic*, *TCP Vegas*, and *LEDBAT never* pass the test, even with a runtime of 60 s. Fig. 4 details the case of *LEDBAT*. The functioning of this congestion-control scheme causes the throughput to ramp-up in the first 38 s and then converge to about 92 Mbps. Thus, if one uses a runtime of 30 s without checking for convergence, the computed mean throughput is about 40 Mbps, which is a totally wrong estimation of *LEDBAT*'s long-running throughput.

**Conclusion.** *TriScale*'s convergence test checks whether the runtime of an experiment is sufficiently long to produce a robust estimate of the long-running performance. A failing convergence test informs a user about the need to increase the runtime or to take other measures (*e.g.*, pruning the start-up time in the raw data) in order to avoid wrong conclusions (*e.g.*, with LEDBAT in Fig. 4).

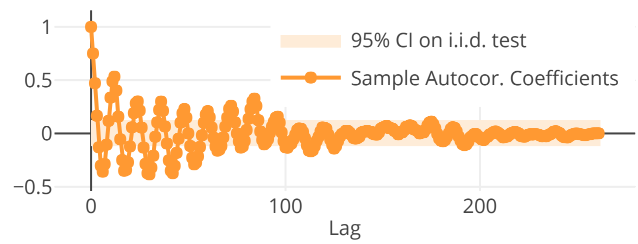


Fig. 5. Autocorrelation plot for the wireless link quality on FlockLab [46], based on the raw data collected by the testbed maintainers [39]. *The dataset contains one test every two hours. The first peak at lag 12 (i.e., 24h) reveals the daily seasonal component. The data also show another at lag 84; which corresponds to one week. Indeed, there is less interference in the weekends than on weekdays, which creates a weekly seasonal component. Data from August 2019. See Appendix § B.2 for further details.*

### 6.2 Wireless Embedded Systems

This case study shows the importance of choosing the time span for a series of runs. In particular, if there are strong temporal patterns in the experimental conditions, one may derive wrong results despite using a high confidence level.

We run a simple evaluation of Glossy [31], a low-power wireless protocol based on synchronous transmissions [83]. A key parameter of Glossy is the number of retransmissions, called $N$. We are hence interested in investigating the impact of two values of $N$ on the reliability of Glossy, measured as the packet reception ratio (PRR). We define our KPI as the median PRR with 95% confidence level (details in § B.2). We collect data using the FlockLab testbed [46], which is located in an office building where we expect more interference during daytime than nighttime. Thus, for each value of $N$, we perform a series of 24 runs scheduled randomly within one day. Computing the KPI leads to a PRR of 88% and 84% for $N = 1$ and $N = 2$, respectively. In other words, it appears that doing two retransmissions instead of one reduces reliability.

*TriScale*'s network profiling function (§ 4.6) provides additional insights. The experiment leads to this (incorrect) conclusion because we neglected a weekly seasonal component, revealed by Fig. 5: there is more interference on weekdays than on weekends. To account for this dependency, we repeat the experiment but extend the span to one week, which leads to KPI values of 80% and 88% for $N$=1 and $N$=2 respectively. This matches our expectation of Glossy's reliability.

**Conclusion.** Using a high confidence level is not enough to avoid drawing wrong conclusions. Real networks exhibit short-term variations that are unpredictable and often unavoidable. This is why it is important to perform multiple runs in a series. Moreover, there may also be systematic patterns, *i.e.*, times with consistently more or less interference. Knowing about and accounting for these patterns is important to ensure fair comparisons: the span of a series should be long enough such that it does not matter when the series starts. *TriScale* helps with those decisions.

### 6.3 Failure Detection

This case study illustrates how the methodology of *TriScale* allows generalizing performance claims for larger sets of input parameters based on a relatively small sample.

We focus on Blink [37], an algorithm that detects failures and reroutes traffic directly in the data plane. The authors evaluated Blink's performance in terms of the true positive rate (TPR, *i.e.*, the fraction of failures successfully detected) and the time needed to reroute the traffic based on 15

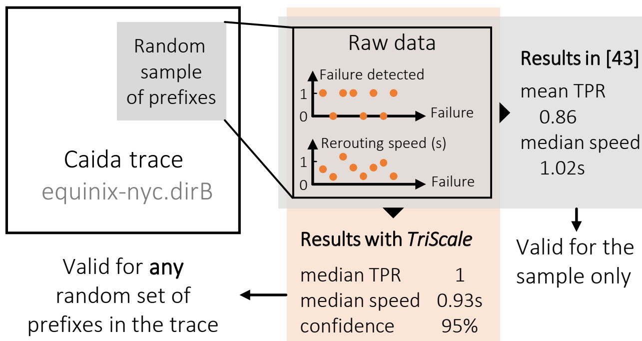


Fig. 6. Using data from a sample of prefixes, *TriScale* allows generalizing and derive performance estimates for any random set of samples from the same Caida trace [22]. *See Appendix § B.3 for further details.*

Internet traces [22, 24] containing data for thousands of prefixes. A subset of prefixes was randomly selected, based on which synthetic traces including artificial failures were generated.

Using *TriScale*, we can generalize the results. For each trace, the evaluation of Blink on one prefix can be seen as a *TriScale* run. Since the prefixes are randomly selected from a fixed set, runs are *i.i.d.* and we can use *TriScale*'s KPI to derive the expected performance of Blink for any set of prefixes (Fig. 6). § B.3 provides more details about Blink's analysis using *TriScale*: We can claim with 95% confidence that, for at least 50% of the prefixes, Blink always detects link failures (TPR= 1) and reroutes traffic within 1 s or less (see Fig. 8).

**Conclusion.** *TriScale*'s methodology can handle any source of performance variability as long as the variability source can be reasonably modeled by a stationary distribution. Thus, one can use *TriScale* to generalize performance claims for evaluations based on network emulation: one can randomly select input traces or system parameters, and derive the expected performance of any other random set. However, the stationarity assumption cannot be always guaranteed (*e.g.*, for cross-traffic over the Internet), which is why *TriScale* includes an empirical independence test.

### 6.4 Video Streaming

This case study shows that the methodology of *TriScale* is compatible with common data reporting practices in networking, such as cumulative density functions (CDF).

In video streaming research, performance is often measured using the quality of experience (QoE) for the user as metric, for example, to compare state-of-the-art adaptive bitrate algorithms such as RobustMPC [82] or Pensieve [48]. Since QoE typically varies a lot, CDFs are often used to give a more global view on the performance of an algorithm. For example, Fig. 7 (shades) shows the CDF achieved by Pensieve over a static set of synthetic network traces (reproduced from [48], see § B.4). However, CDFs are no different from other metrics: What is the confidence in the result? How much would it vary with a different set of traces?

A CDF is a representation of all percentiles of a given distribution. Hence, *TriScale* can be directly used to estimate an *entire CDF* by computing a large set of KPIs. For example, Fig. 7 (solid line) shows the 95% CI for the 2th to the 98th percentile, which provides a lower-bound on the expected performance. Hence, one can claim that, for *any* set of traces that would be generated/obtained similarly, the QoE of an algorithm is better than the CI CDF with 95% confidence.

**Conclusion.** Thanks to the use of percentiles as KPI, *TriScale* is agnostic to the choice of metric and handles any source of variability that can be modeled as a stationary distribution.

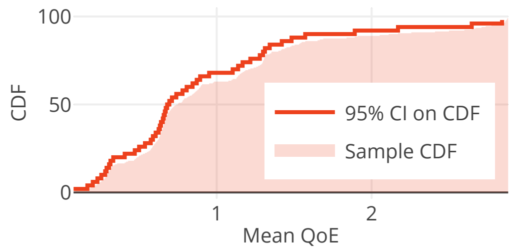


Fig. 7. A CDF and its 95% CI, computed by *TriScale. Original CDF reproduced from [48]. The CI provides a lower-bound on the expected performance for any other random set of input traces generated similarly. See Appendix § B.4 for details.*

## 7 DISCUSSION AND FUTURE WORK

**Data collection.** *TriScale* is not responsible for the execution of networking experiments: it does not perform the data collection. Frameworks such as Pantheon [81] or Puffer [80] are specialized in data collection; other examples include low-power wireless testbeds [46, 65, 66] and networking facilities [11, 29, 55]. *TriScale* can be integrated into these frameworks to create a fully-automated experimentation chain and build full-fledged benchmarking infrastructures [16].

**Human-in-the-loop.** *TriScale* automates the data analysis and implements tests that verify whether the required hypotheses hold. It is up to the user to critically assess why tests fail when they do (*e.g.*, a longer runtime needed for *LEDBAT* – § 6.1), and derive corresponding countermeasure (*e.g.*, pruning the start-up time in the raw data or re-running experiments). Furthermore, there will likely be some feedback and iterations between the first set of tests and the final evaluation as a larger set of runs may uncover insights such as unknown correlation or seasonal components in the data.

**Ranking solutions.** *TriScale* compares performance, but it does not rank. The evaluation results are always relative to a specific network or evaluation scenario (*e.g.*, a given cloud provider [72]). It is not trivial to generalize and claim that a solution A is better than a solution B. This problem relates to benchmarking and multi-objective optimization, which goes beyond the scope of *TriScale*.

**Community guidelines.** *TriScale* formalizes evaluation objectives (§ 4.4), but it does not dictate which parameters to use. Similarly, *TriScale* quantifies the replicability of an experiment (§ 4.3), but it does not conclude whether the evaluation is replicable or not (§ 4.7). It is now up to the networking communities to set their own standards, reference parameters to use, and acceptable requirements, as already done in other disciplines [33].

**Limitations.** As mentioned throughout the paper, *TriScale* is a *systematic methodology* that streamlines the design and analysis of performance evaluations and that supports the replicability of networking evaluations. The specific models and methods we present here are tailored towards a large class of performance evaluation scenarios, exemplified by the various case studies in § 6. Their usefulness and expressiveness is based on available or tested previous knowledge about the system under evaluation. Other systems may have different behavior, which may lead to other models and methods being more appropriate to apply at the different stages of the methodology. Relevant examples include the definition of convergence (which may be different depending on the system), normalization of measurements using different classes of scaling function, the sampling of runs within series and series within sequels (period vs. random vs. biased random), or available knowledge about the distributions and other statistical properties of measurements.

## 8 RELATED WORK

The replicability of experiments and comparability of results are cornerstones of the scientific method. In recent years, several studies have highlighted the inability of researchers from various disciplines to replicate their own experimental results [10, 58], often due to sloppy research protocols and faulty statistical analysis [15, 17, 64]. This is a problem in computer science as well [26, 73], where experiments are seldom replicable and artifacts rarely shared.

**Promoting replicability.** Recent work demonstrated that poor experimental and statistical practices has led to wrong or ambiguous conclusions. [72] presents a survey of recent cloud computing works and concludes that more than 60% of papers reports poor or no specification of the experiments and three-quarter of those that do use less repetition than necessary to mitigate the performance variability of cloud infrastructure. [80] shows that even with 2.5 years of data, the 95% CI of some performance metrics of adaptive bit-rate algorithms are of the same size as the performance "improvements" claimed in the original papers. To address this "replicability crisis" [10], many efforts aiming to incentivize a rigorous experimentation have gained momentum in computer science, including *e.g.*, ACM's badging system for publications [1]. Especially in the networking community – challenged by the need to carry out experiments on dynamic and uncontrollable conditions [21, 51] – several workshops [8, 19, 35], surveys [32], and guidelines [9, 43, 52, 63] have raised awareness on the replicability problem and promoted better experimentation practices. This large body of work mostly offers *qualitative* statements on how an experiment should be performed and documented. Such statements emphasize, *e.g.*, the need to carefully choose when and how often to sample data [9], or suggest which methodology to adopt during performance evaluations [43]. However, there is no guarantee that following these recommendations leads to replicable results, nor is there a concrete way to assess whether an experiment can be considered replicable.

In contrast, *TriScale* provides researchers with *quantitative* answers about how to concretely perform an experiment, *e.g.*, how many runs should be performed and how long should they be; answers derived by following a clear experimental methodology grounded on robust non-parametric statistics. Moreover, *TriScale* offers a way to assess and compare the replicability of experimental results using unambiguous performance indicators and variability scores.

**Supporting replicability.** A large number of experimental facilities and tools have been developed to aid researchers in carrying out replicable networking studies [55]. Testbeds such as EmuLab [75] and FlexLab [61], as well as emulation tools such as MiniNet [36], enable the creation of artificial network conditions using a given specification or passively-observed traffic. Emulated conditions offer a more controlled environment than experiments with real-world traffic (*e.g.*, by transmitting data over the Internet [14, 25], cloud [18, 29], or wireless interfaces [2, 34, 50]). Still, they suffer from performance variability caused by the underlying hardware and software components, which hampers replicability [49]. To overcome these problems, several solutions have been proposed [30]: *e.g.*, revisiting OS libraries [69], using virtualization [36, 41, 42], adaptable profiles [62], and fault patterns [3]. Other tools have been developed to support mobility experiments [11, 23], maximize the repeatability of interference generation [67], and enable researchers to consistently evaluate congestion-control schemes [81].

All the aforementioned tools aim to improve replicability *during* the experiments, while *TriScale* assists researchers *before* and *after* their execution. It does so by informing about the number and length of runs necessary to reach a given level of confidence, as well as by computing a score quantifying the variability of the results. Hence, *TriScale* complements the existing body of literature promoting and enhancing replicability in networking research. Another tool similar to *TriScale* is CONFIRM [49]; it uses the same method to compute confidence intervals but aims to predict the number of runs necessary to obtain intervals of a given size, *e.g.*, ±1% of the median.

Such predictions are necessarily context-specific (cloud computing in this case), whereas *TriScale* approach is more general; it indicates how many samples are required to compute a CI but does not say anything about the interval size.

In a prior workshop paper [4], we have argued that a well-defined methodology specifying how to plan, execute, and report on experimental results is of paramount importance for the networking community: *TriScale* is the concrete realization of this early vision into a tangible framework.

## 9 CONCLUSIONS

A consistent methodology for the design and analysis of networking experiments is crucial for a more rigorous and replicable scientific activity. *TriScale* implements a methodology grounded on non-parametric statistics into a framework that aids researchers in designing experiments and analyzing data. *TriScale* ultimately improves the interpretation of results and helps to quantify the replicability of experimental evaluations.

We hope that *TriScale*'s open availability and usability [5, 6] will foster better experimentation practices in the short term for the networking community at large. The quest towards fully-replicable networking experiments remains open, but we believe *TriScale* represents a first stepping stone towards an accepted standard for experimental evaluations.

## REFERENCES


[1] ACM. 2018. Artifact Review and Badging. https://www.acm.org/publications/policies/artifact-review-badging.

[2] Cedric Adjih, Emmanuel Baccelli, Eric Fleury, Gaetan Harter, Nathalie Mitton, Thomas Noel, Roger Pissard-Gibollet, Frederic Saint-Marcel, Guillaume Schreiner, Julien Vandaele, and Thomas Watteyne. 2015. FIT IoT-LAB: A Large Scale Open Experimental IoT Testbed. In *2015 IEEE 2nd World Forum on Internet of Things (WF-IoT)*. https://doi.org/10.1109/WF-IoT.2015.7389098

[3] Angainor. [n.d.]. Angainor – Reproducible Evaluation and Fault Injection of Large-Scale Distributed Systems. http://angainor.science/.

[4] Anonymous. [n.d.]. Reference Blinded Due to Double-Blind Submission.

[5] Anonymous. [n.d.]. TriScale-Anon/Triscale. https://github.com/TriScale-Anon/triscale.

[6] Anonymous. [n.d.]. TriScale Live Demo. https://mybinder.org/v2/gh/TriScale-Anon/triscale/master.

[7] Anonymous. [n.d.]. TriScale's Artifact Repository. https://doi.org/10.5281/zenodo.3451417.

[8] Vaibhav Bajpai, Olivier Bonaventure, Kimberly Claffy, and Daniel Karrenberg. 2019. Encouraging Reproducibility in Scientific Research of the Internet (Dagstuhl Seminar 18412). *Dagstuhl Reports* (2019). https://doi.org/10.4230/DagRep.8.10.41

[9] Vaibhav Bajpai, Anna Brunstrom, Anja Feldmann, Wolfgang Kellerer, Aiko Pras, Henning Schulzrinne, Georgios Smaragdakis, Matthias Wählisch, and Klaus Wehrle. 2019. The Dagstuhl Beginners Guide to Reproducibility for Experimental Networking Research. *SIGCOMM Comput. Commun. Rev.* (2019). https://dl.acm.org/citation.cfm?id=3314217.

[10] Monya Baker. 2016. Is There a Reproducibility Crisis? *Nature News* (2016). https://www.nature.com/news/polopoly_fs/1.19970!/menu/main/topColumns/topLeftColumn/pdf/533452a.pdf.

[11] Arijit Banerjee, Junguk Cho, Eric Eide, Jonathon Duerig, Binh Nguyen, Robert Ricci, Jacobus Van der Merwe, Kirk Webb, and Gary Wong. 2015. PhantomNet: Research Infrastructure for Mobile Networking, Cloud Computing and Software-Defined Networking. *GetMobile: Mobile Comp. and Comm.* (2015). https://doi.org/10.1145/2817761.2817772

[12] Lorena A. Barba. 2018. Terminologies for Reproducible Research. *arXiv:1802.03311 [cs]* (2018). arXiv:1802.03311 [cs] http://arxiv.org/abs/1802.03311.

[13] Stephanie R. Barbari, Daniel P. Kane, Elizabeth A. Moore, and Polina V. Shcherbakova. 2018. Functional Analysis of Cancer-Associated DNA Polymerase $\epsilon$ Variants in Saccharomyces Cerevisiae. *G3: Genes, Genomes, Genetics* (2018). https://doi.org/10.1534/g3.118.200042

[14] Mark Berman, Jeffrey S. Chase, Lawrence Landweber, Akihiro Nakao, Max Ott, Dipankar Raychaudhuri, Robert Ricci, and Ivan Seskar. 2014. GENI: A Federated Testbed for Innovative Network Experiments. *Computer Networks* (2014). https://doi.org/10.1016/j.bjp.2013.12.037

[15] Stephen M. Blackburn, Amer Diwan, Matthias Hauswirth, Peter F. Sweeney, José Nelson Amaral, Tim Brecht, Lubomír Bulej, Cliff Click, Lieven Eeckhout, Sebastian Fischmeister, Daniel Frampton, Laurie J. Hendren, Michael Hind, Antony L. Hosking, Richard E. Jones, Tomas Kalibera, Nathan Keynes, Nathaniel Nystrom, and Andreas Zeller. 2016. The Truth,

The Whole Truth, and Nothing But the Truth: A Pragmatic Guide to Assessing Empirical Evaluations. *ACM Trans. Program. Lang. Syst.* (2016). https://doi.org/10.1145/2983574
[16] Carlo A. Boano, Simon Duquennoy, Anna Förster, Omprakash Gnawali, Romain Jacob, Hyung-Sin Kim, Olaf Landsiedel, Ramona Marfievici, Luca Mottola, Gian Pietro Picco, Xavier Vilajosana, Thomas Watteyne, and Marco Zimmerling. 2018. IoTBench: Towards a Benchmark for Low-Power Wireless Networking. In *1st Workshop on Benchmarking Cyber-Physical Networks and Systems (CPSBench 2018).* https://doi.org/10.3929/ethz-b-000256517
[17] Ronald F. Boisvert. 2016. Incentivizing Reproducibility. *Commun. ACM* (2016). https://doi.org/10.1145/2994031
[18] Raphaël Bolze et al. 2006. Grid'5000: A Large Scale And Highly Reconfigurable Experimental Grid Testbed. *International Journal of High Performance Computing Applications* (2006). https://hal.inria.fr/hal-00684943.
[19] Olivier Bonaventure, Luigi Iannone, and Damien Saucez (Eds.). 2017. *Proceedings of the International ACM SIGCOMM Reproducibility Workshop.* ACM, Los Angeles, CA, USA. https://dl.acm.org/citation.cfm?id=3097766.
[20] Peter J. Brockwell, Richard A. Davis, and Stephen E. Fienberg. 1991. *Time Series: Theory and Methods: Theory and Methods.* Springer Science & Business Media. https://doi.org/10.1007/978-1-4419-0320-4
[21] Ryan Burchfield, Ehsan Nourbakhsh, Jeff Dix, Kunal Sahu, S. Venkatesan, and Ravi Prakash. 2009. RF in the Jungle: Effect of Environment Assumptions on Wireless Experiment Repeatability. In *Proceedings of the International Conference on Communications (ICC).* IEEE. https://ehsaan.net/wp-content/uploads/publications/rfij.pdf.
[22] CAIDA. [n.d.]. CAIDA Internet Data – Passive Data Sources. https://www.caida.org/data/passive/index.xml.
[23] Junguk Cho, Jonathan Duerig, Eric Eide, Binh Nguyen, Robert Ricci, Aisha Syed, Jacobus Van der Merwe, Kirk Webb, and Gary Wong. 2016. Repeatable Mobile Networking Research with phantomNet: Demo. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking - MobiCom '16.* ACM Press, New York City, New York. https://doi.org/10.1145/2973750.2985616
[24] Kenjiro Cho, Koushirou Mitsuya, and Akira Kato. 2000. Traffic Data Repository at the WIDE Project. In *Proceedings of the Annual Conference on USENIX Annual Technical Conference (ATEC '00).* USENIX Association, San Diego, California.
[25] Brent Chun, David Culler, Timothy Roscoe, Andy Bavier, Larry Peterson, Mike Wawrzoniak, and Mic Bowman. 2003. PlanetLab: An Overlay Testbed for Broad-Coverage Services. *SIGCOMM Comput. Commun. Rev.* (2003). https://doi.org/10.1145/956993.956995
[26] Christian Collberg, Todd Proebsting, and Alex M. Warren. 2015. *Repeatability and Benefaction in Computer Systems Research.* Technical Report. University of Arizona. http://reproducibility.cs.arizona.edu/v2/RepeatabilityTR.pdf.
[27] Geoff Cumming and Sue Finch. 2001. A Primer on the Understanding, Use, and Calculation of Confidence Intervals That Are Based on Central and Noncentral Distributions. *Educational and Psychological Measurement* (2001). https://doi.org/10.1177/0013164401614002.
[28] Herbert A. David and Haikady N. Nagaraja. 2005. Order Statistics in Nonparametric Inference. In *Order Statistics.* John Wiley & Sons, Ltd. https://doi.org/10.1002/0471722162.ch7
[29] Dmitry Duplyakin, Robert Ricci, Aleksander Maricq, Gary Wong, Jonathon Duerig, Eric Eide, Leigh Stoller, Mike Hibler, David Johnson, Kirk Webb, Aditya Akella, Kuangching Wang, Glenn Ricart, Larry Landweber, Chip Elliott, Michael Zink, Emmanuel Cecchet, Snigdhaswin Kar, and Prabodh Mishra. 2019. The Design and Operation of CloudLab. In *2019 USENIX Annual Technical Conference (USENIX ATC 19).* https://www.usenix.org/conference/atc19/presentation/duplyakin.
[30] Sarah Edwards, Xuan Liu, and Niky Riga. 2015. Creating Repeatable Computer Science and Networking Experiments on Shared, Public Testbeds. *SIGOPS Oper. Syst. Rev.* (2015). http://doi.acm.org/10.1145/2723872.2723884.
[31] Federico Ferrari, Marco Zimmerling, Lothar Thiele, and Olga Saukh. 2011. Efficient Network Flooding and Time Synchronization with Glossy. In *Proceedings of the 10th ACM/IEEE International Conference on Information Processing in Sensor Networks.* https://ieeexplore.ieee.org/document/5779066.
[32] Matthias Flittner, Mohamed Naoufal Mahfoudi, Damien Saucez, Matthias Wählisch, Luigi Iannone, Vaibhav Bajpai, and Alex Afanasyev. 2018. A Survey on Artifacts from CoNEXT, ICN, IMC, and SIGCOMM Conferences in 2017. *SIGCOMM Comput. Commun. Rev.* (2018). https://doi.org/10.1145/3211852.3211864
[33] C. Galán, Matt Smith, M. Thibaudon, G. Frenguelli, J. Oteros, R. Gehrig, U. Berger, B. Clot, R. Brandao, and EAS QC Working Group. 2014. Pollen Monitoring: Minimum Requirements and Reproducibility of Analysis. *Aerobiologia* (2014). https://doi.org/10.1007/s10453-014-9335-5
[34] Sachin Ganu, Haris Kremo, Richard Howard, and Ivan Seskar. 2005. Addressing Repeatability in Wireless Experiments Using ORBIT Testbed. In *Proceedings of the 1st International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TRIDENTCOM).* IEEE Computer Society, Trento, Italy. https://ieeexplore.ieee.org/document/1386191.
[35] Omprakash Gnawali, Marco Zimmerling, and Sebastian Trimpe (Eds.). 2018. *Proceedings of the 1st International Workshop on Benchmarking Cyber-Physical Networks and Systems (CPSBench).* IEEE, Porto, Portugal. https://doi.org/10.1109/CPSBench.2018.00004.

[36] Nikhil Handigol, Brandon Heller, Vimalkumar Jeyakumar, Bob Lantz, and Nick McKeown. 2012. Reproducible Network Experiments Using Container-Based Emulation. In *Proceedings of the 8th International Conference on Emerging Networking Experiments and Technologies (CoNEXT)*. ACM, Nice, France. http://tiny-tera.stanford.edu/~nickm/papers/p253.pdf.

[37] Thomas Holterbach, Edgar Costa Molero, Maria Apostolaki, Alberto Dainotti, Stefano Vissicchio, and Laurent Vanbever. 2019. Blink: Fast Connectivity Recovery Entirely in the Data Plane. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*. https://www.usenix.org/conference/nsdi19/presentation/holterbach.

[38] Romain Jacob, Carlo Alberto Boano, Usman Raza, Marco Zimmerling, and Lothar Thiele. 2019. Towards a Methodology for Experimental Evaluation in Low-Power Wireless Networking. In *2nd Workshop on Benchmarking Cyber-Physical Systems and Internet of Things (CPS-IoTBench'19)*. https://doi.org/10.3929/ethz-b-000325096

[39] Romain Jacob, Reto Da Forno, Roman Trüb, Andreas Biri, and Lothar Thiele. 2019. Dataset: Wireless Link Quality Estimation on FlockLab – and Beyond. In *Proceedings of the 2nd International Workshop on Data Acquisition to Analysis (DATA)*. ACM, New York, NY, USA. https://doi.org/10.3929/ethz-b-000355846

[40] Project Jupyter, Matthias Bussonnier, Jessica Forde, Jeremy Freeman, Brian Granger, Tim Head, Chris Holdgraf, Kyle Kelley, Gladys Nalvarte, Andrew Osheroff, M. Pacer, Yuvi Panda, Fernando Perez, Benjamin Ragan-Kelley, and Carol Willing. 2018. Binder 2.0 - Reproducible, Interactive, Sharable Environments for Science at Scale. In *Proceedings of the 17th Python in Science Conference*. https://doi.org/10.25080/Majora-4af1f417-011

[41] Pravein Govindan Kannan, Ahmad Soltani, Mun Choon Chan, and Ee-Chien Chang. 2018. BNV: Enabling Scalable Network Experimentation throughBare-Metal Network Virtualization. In *Proceedings of the 11th USENIX Conference on Cyber Security Experimentation and Test (CSET)*. USENIX Association. http://dl.acm.org/citation.cfm?id=3307412.3307418.

[42] Teemu Koponen, Keith Amidon, Peter Balland, Martín Casado, Anupam Chanda, Bryan Fulton, Igor Ganichev, Jesse Gross, Paul Ingram, Ethan Jackson, Andrew Lambeth, Romain Lenglet, Shih-Hao Li, Amar Padmanabhan, Justin Pettit, Ben Pfaff, Rajiv Ramanathan, Scott Shenker, Alan Shieh, Jeremy Stribling, Pankaj Thakkar, Dan Wendlandt, Alexander Yip, and Ronghua Zhang. 2014. Network Virtualization in Multi-Tenant Datacenters. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. https://www.usenix.org/conference/nsdi14/technical-sessions/presentation/koponen.

[43] K. Kritsis, G. Z. Papadopoulos, A. Gallais, P. Chatzimisios, and F. Théoleyre. 2018. A Tutorial on Performance Evaluation and Validation Methodology for Low-Power and Lossy Networks. *IEEE Communications Surveys Tutorials* (2018). https://doi.org/10.1109/COMST.2018.2820810

[44] Daniël Lakens. 2017. Equivalence Tests: A Practical Primer for *t* Tests, Correlations, and Meta-Analyses. *Social Psychological and Personality Science* (2017). https://doi.org/10.1177/1948550617697177

[45] Jean-Yves Le Boudec. 2011. *Performance Evaluation of Computer and Communication Systems*. Epfl Press.

[46] Roman Lim, Federico Ferrari, Marco Zimmerling, Christoph Walser, Philipp Sommer, and Jan Beutel. 2013. FlockLab: A Testbed for Distributed, Synchronized Tracing and Profiling of Wireless Embedded Systems. In *Proceedings of the 12th International Conference on Information Processing in Sensor Networks (IPSN'13)*. ACM, New York, NY, USA. https://doi.org/10.1145/2461381.2461402

[47] David M. Liu and Matthew J. Salganik. 2019. *Successes and Struggles with Computational Reproducibility: Lessons from the Fragile Families Challenge*. Technical Report. OSF.io. https://osf.io/preprints/socarxiv/g3pdb/.

[48] Hongzi Mao, Ravi Netravali, and Mohammad Alizadeh. 2017. Neural Adaptive Video Streaming with Pensieve. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication (SIGCOMM'17)*. Association for Computing Machinery, Los Angeles, CA, USA. https://doi.org/10.1145/3098822.3098843

[49] Aleksander Maricq, Dmitry Duplyakin, Ivo Jimenez, Carlos Maltzahn, Ryan Stutsman, and Robert Ricci. 2018. Taming Performance Variability. In *Proceedings of the 13th International USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. USENIX Association, Carlsbad, CA, USA. https://www.usenix.org/system/files/osdi18-maricq.pdf.

[50] Abdelbassat Massouri, Leonardo Cardoso, Benjamin Guillon, Florin Hutu, Guillaume Villemaud, Tanguy Risset, and Jean-Marie Gorce. 2014. CorteXlab: An Open FPGA-Based Facility for Testing SDR and Cognitive Radio Networks in a Reproducible Environment. In *Proceedings of the International Conference on Computer Communications (INFOCOM) Workshops*. IEEE, San Francisco, CA, USA. https://ieeexplore.ieee.org/document/6849176.

[51] Miguel Matos. 2018. Towards Reproducible Evaluation of Large-Scale Distributed Systems. In *Proceedings of the International Workshop on Advanced Tools, Programming Languages, and Platforms for Implementing and Evaluating Algorithms for Distributed Systems (ApPLIED)*. ACM, Egham, United Kingdom. https://doi.org/10.1145/3231104.3231113

[52] Micro Focus. 2018. *Seven Ways to Fail*. Technical Report. –. https://www.microfocus.com/media/brochure/seven_ways_to_fail_brochure.pdf.

[53] Ravi Netravali, Anirudh Sivaraman, Somak Das, Ameesh Goyal, Keith Winstein, James Mickens, and Hari Balakrishnan. 2015. Mahimahi: Accurate Record-and-Replay for HTTP. In *Proceedings of the International USENIX Annual Technical*

*Conference (ATC)*. USENIX Association, Santa Clara, CA, USA. https://www.usenix.org/system/files/conference/atc15/atc15-paper-netravali.pdf.
[54] Numpy. [n.d.]. NumPy: The Fundamental Package for Scientific Computing with Python. https://numpy.org/.
[55] Lucas Nussbaum. 2017. Testbeds Support for Reproducible Research. In *Proceedings of the International ACM SIGCOMM Reproducibility Workshop*. ACM, Los Angeles, CA, USA. https://doi.org/10.1145/3097766.3097773
[56] Pantheon. [n.d.]. Pantheon. https://pantheon.stanford.edu/.
[57] Vern Paxson. 2004. Strategies for Sound Internet Measurement. In *Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement (IMC '04)*. Association for Computing Machinery, Taormina, Sicily, Italy. https://doi.org/10.1145/1028788.1028824
[58] Roger Peng. 2015. The Reproducibility Crisis in Science: A Statistical Counterattack. *Significance* (2015). https://doi.org/10.1111/j.1740-9713.2015.00827.x.
[59] Hans E. Plesser. 2018. Reproducibility vs. Replicability: A Brief History of a Confused Terminology. *Frontiers in Neuroinformatics* (2018). https://doi.org/10.1177/1948550617697177.
[60] Plotly. [n.d.]. Plotly: Modern Analytic Apps for the Enterprise. https://plot.ly.
[61] Robert Ricci, Jonathon Duerig, Pramod Sanaga, Daniel Gebhardt, Mike Hibler, Kevin Atkinson, Junxing Zhang, Sneha Kasera, and Jay Lepreau. 2007. The Flexlab Approach to Realistic Evaluation of Networked Systems. In *Proceedings of the 4th USENIX Conference on Networked Systems Design & Implementation (NSDI'07)*. USENIX Association, Cambridge, MA, USA. https://www.cs.utah.edu/flux/papers/flexlab-nsdi07.pdf.
[62] Robert Ricci, Gary Wong, Leigh Stoller, Kirk Webb, Jonathon Duerig, Keith Downie, and Mike Hibler. 2015. Apt: A Platform for Repeatable Research in Computer Science. *SIGOPS Oper. Syst. Rev.* (2015). https://doi.org/10.1145/2723872.2723885
[63] Damien Saucez and Luigi Iannone. 2018. Thoughts and Recommendations from the ACM SIGCOMM 2017 Reproducibility Workshop. *SIGCOMM Comput. Commun. Rev.* (2018). https://doi.org/10.1145/3211852.3211863
[64] Hanspeter Schmid and Alex Huber. 2014. Measuring a Small Number of Samples, and the $3\sigma$ Fallacy: Shedding Light on Confidence and Error Intervals. *IEEE Solid-State Circuits Magazine* (2014). https://doi.org/10.1109/MSSC.2014.2313714
[65] Markus Schuß, Carlo Alberto Boano, and Kay Römer. 2018. Moving Beyond Competitions: Extending D-Cube to Seamlessly Benchmark Low-Power Wireless Systems. In *Proceedings of the 1st International Workshop on Benchmarking Cyber-Physical Networks and Systems (CPSBench)*. IEEE. https://doi.org/10.1109/CPSBench.2018.00012
[66] Markus Schuß, Carlo Alberto Boano, Manuel Weber, and Kay Römer. 2017. A Competition to Push the Dependability of Low-Power Wireless Protocols to the Edge. In *Proceedings of the 2017 International Conference on Embedded Wireless Systems and Networks (EWSN'17)*. Junction Publishing, USA. https://doi.org/10.5555/3108009.3108018
[67] Markus Schuß, Carlo Alberto Boano, Manuel Weber, Matthias Schulz, Matthias Hollick, and Kay Römer. 2019. JamLab-NG: Benchmarking Low-Power Wireless Protocols under Controllable and Repeatable Wi-Fi Interference. In *Proceedings of the 2019 International Conference on Embedded Wireless Systems and Networks (EWSN '19)*. Junction Publishing, Beijing, China. https://dl.acm.org/doi/abs/10.5555/3324320.3324331.
[68] Pranab Kumar Sen. 1968. Estimates of the Regression Coefficient Based on Kendall's Tau. *J. Amer. Statist. Assoc.* (1968). https://doi.org/10.1080/01621459.1968.10480934
[69] Hajime Tazaki, Frédéric Uarbani, Emilio Mancini, Mathieu Lacage, Daniel Camara, Thierry Turletti, and Walid Dabbous. 2013. Direct Code Execution: Revisiting Library OS Architecture for Reproducible Network Experiments. In *Proceedings of the 9th International Conference on Emerging Networking Experiments and Technologies (CoNEXT) (CoNEXT'13)*. ACM, New York, NY, USA. https://doi.org/10.1145/2535372.2535374
[70] Henri Theil. 1992. A Rank-Invariant Method of Linear and Polynomial Regression Analysis. In *Henri Theil's Contributions to Economics and Econometrics: Econometric Theory and Methodology*, Baldev Raj and Johan Koerts (Eds.). Springer Netherlands, Dordrecht. https://doi.org/10.1007/978-94-011-2546-8_20
[71] William R. Thompson. 1936. On Confidence Ranges for the Median and Other Expectation Distributions for Populations of Unknown Distribution Form. *The Annals of Mathematical Statistics* (1936). https://www.jstor.org/stable/2957563.
[72] Alexandru Uta, Alexandru Custura, Dmitry Duplyakin, Ivo Jimenez, Jan Rellermeyer, Carlos Maltzahn, Robert Ricci, and Alexandru Iosup. 2020. Is Big Data Performance Reproducible in Modern Cloud Networks?. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. https://www.usenix.org/conference/nsdi20/presentation/uta.
[73] Jan Vitek and Tomas Kalibera. 2011. Repeatability, Reproducibility and Rigor in Systems Research. In *Proceedings of the 9th International Conference on Embedded Software (EMSOFT)*. ACM. https://www.cs.kent.ac.uk/pubs/2011/3174/content.pdf.
[74] Ronald L. Wasserstein, Allen L. Schirm, and Nicole A. Lazar. 2019. Moving to a World Beyond "$p < 0.05$". *The American Statistician* (2019). https://doi.org/10.1080/00031305.2019.1583913
[75] Brian White, Jay Lepreau, Leigh Stoller, Robert Ricci, Shashi Guruprasad, Mac Newbold, Mike Hibler, Chad Barb, and Abhijeet Joglekar. 2002. An Integrated Experimental Environment for Distributed Systems and Networks. *SIGOPS*

*Oper. Syst. Rev.* (2002). http://doi.acm.org/10.1145/844128.844152.

[76] Wikipedia. 2019. Kruskal–Wallis One-Way Analysis of Variance. *Wikipedia* (2019). https://en.wikipedia.org/w/index.php?title=Kruskal%E2%80%93Wallis_one-way_analysis_of_variance&oldid=930537699.

[77] Wikipedia. 2019. One-Way Analysis of Variance. *Wikipedia* (2019). https://en.wikipedia.org/w/index.php?title=One-way_analysis_of_variance&oldid=931491661.

[78] Wikipedia. 2019. Theil–Sen Estimator. *Wikipedia* (2019). https://en.wikipedia.org/w/index.php?title=Theil%E2%80%93Sen_estimator&oldid=906073990.

[79] Wikipedia. 2020. Central Limit Theorem. *Wikipedia* (2020). https://en.wikipedia.org/w/index.php?title=Central_limit_theorem&oldid=937666343.

[80] Francis Y. Yan, Hudson Ayers, Chenzhi Zhu, Sadjad Fouladi, James Hong, Keyi Zhang, Philip Levis, and Keith Winstein. 2020. Learning in Situ: A Randomized Experiment in Video Streaming. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. https://www.usenix.org/conference/nsdi20/presentation/yan.

[81] Francis Y. Yan, Jestin Ma, Greg D. Hill, Deepti Raghavan, Riad S. Wahby, Philip Levis, and Keith Winstein. 2018. Pantheon: The Training Ground for Internet Congestion-Control Research. In *Proceedings of the International USENIX Annual Technical Conference (ATC)*. USENIX Association, Boston, MA, USA. https://www.usenix.org/conference/atc18/presentation/yan-francis.

[82] Xiaoqi Yin, Abhishek Jindal, Vyas Sekar, and Bruno Sinopoli. 2015. A Control-Theoretic Approach for Dynamic Adaptive Video Streaming over HTTP. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication (SIGCOMM '15)*. Association for Computing Machinery, London, United Kingdom. https://doi.org/10.1145/2785956.2787486

[83] Marco Zimmerling, Luca Mottola, and Silvia Santini. 2020. Synchronous Transmissions in Low-Power Wireless: A Survey of Communication Protocols and Network Services. *arXiv:2001.08557 [cs, eess]* (2020). arXiv:2001.08557 [cs, eess]

## A DETAILS ON THE SCALABILITY EVALUATION

This appendix provides additional information about the evaluation of *TriScale*'s scalability presented in § 5.2. We perform the evaluation using a Jupyter notebook[5] (*i.e.*, an open-source web-based interactive computational environment to create and share documents containing live code, equations, visualizations, and text) that is available in the *TriScale* repository [5]. Such evaluation, which we run on a commodity laptop, yields the results summarized in Table 2.

**Results – Metrics.** The data shows two modes in the execution time of the *analysis_metric*() function: a step increase, followed by a slow linear increase. This can be easily explained: the more computationally expensive part of *analysis_metric*() is the convergence test, which includes the Theil-Sen regression (§ 4.5). The latter works by computing the slopes between all pairs of points and returns the median slope value; thus, the regression scales with $O(n^2)$.

However, *TriScale* does not perform the regression on the input data directly. Instead, *TriScale* divides the input data in chunks. For each chunk, a metric value is computed, leading to a new data series of metric values. The purpose of the convergence test is to verify that these metric values have converged; thus *TriScale* executes the Theil-Sen regression on this new data series. The Theil-Sen regression does not require many samples for producing a reliable result; a few tens of data points are often considered sufficient [78]. Thus, we can cap the size of metric data series (*TriScale* caps it to 100 values – § 4.1), which bounds the execution time of the Theil-Sen regression. Ultimately, this allows the *analysis_metric*() function to scale very well with the sample size.

The linear increase for a large number of raw samples is due to the computation of the metric on increasingly large chunks. The more complex the metric is, the longer the execution time. In this evaluation, a percentile is used as metric, which is computed efficiently with NumPy [54].

Overall, running *analysis_metric*() takes about 1 s for up to one million data points. The data collection time depends on the networking experiment, but it is unlikely that many experiments would produce much more than a million of data points per second. Thus, we conclude that the computation time of the *analysis_metric*() function is negligible for networking experiments.

**Results – KPIs.** The data shows a clear linear correlation between the sample size and the execution time of the *analysis_kpi*() function, which is not surprising: most computations are related to the determination of the confidence interval using Thompson's method, which is an iterative process through the ordered data samples [71].

The input size for the KPI computation is the number of series one performs for an experiment. Our results show it takes less than 100 ms for an input size of 1000; we thus conclude that the computation time of the *analysis_kpi*() function is negligible for networking experiments.

**Results - Variability scores.** Unsurprisingly, the data is very similar as for *analysis_kpi*(): The two functions essentially perform the same computations. They only differ in the generation of outputs (logs and plots). Since the outputs are not considered in this scalability evaluation, we obtain very similar results for both functions. Thus, we conclude that the computation time of the *analysis_variability*() function is negligible for networking experiments.

## B DETAILS ON CASE STUDIES

This appendix provides details on the four case studies presented in § 6; in particular, it details each evaluation scenario and how we have obtained the data. All case studies are performed using Jupyter notebooks, which are available in the *TriScale* repository [5].

[5]triscale_scalability.ipynb

Table 2. Scalability evaluation. *TriScale data analysis is fast and scales well with increasing input sizes. The most time-consuming element is the convergence test (§ 4.5) which is performed before the computation of metrics. Still, it generally takes less than one second for inputs (i.e., the number of raw measurements in a run) of up to one million data points.*

| **Computation of** | **Input size** | **Execution time** (approx.) |
|---|---|---|
| Metrics | 1000 | 20 ms |
| | 10 k | 50 ms |
| | 1 M | 1 s |
| KPIs and Variability scores | 100 | 10 ms |
| | 1000 | 100 ms |

### B.1 Congestion Control

**Reproducing the case study.** The entire case study is described in detail in a Jupyter notebook[6] which is available in the *TriScale* repository [5].

**Evaluation scenario.** This case study compares the performance of 17 congestion-control schemes using Pantheon [81]. We evaluate the throughput and one-way delay of full-throttle flows, *i.e.*, stable flows whose only throttling/limiting factor is the congestion control. For a fair comparison between the schemes, we use the MahiMahi emulator [53] (integrated in Pantheon) and focus on a single flow scenario. We use only the calibrated path from AWS California to Mexico, provided by Pantheon.[7]

**Data collection.** We build the Pantheon toolchain from the source code provided by the authors[8] and test all schemes locally based on the aforementioned emulated network. We only modify the authors' code to save the throughput and delay raw data, such that we can do the analysis of runs using *TriScale*. We perform two sets of experiments with always 10 runs per series:

- A set of 5 series with a runtime of 30 s.
- A set of series with a runtime of 10, 20, 40, 50, and 60 s, respectively (one of each).

The data we collected are available on Zenodo [7].

### B.2 Wireless Embedded Systems

**Reproducing the case study.** The entire case study is described in detail in a Jupyter notebook[9] which is available in the *TriScale* repository [5].

**Evaluation scenario.** We run a simple evaluation of Glossy [31], a low-power wireless protocol which includes as parameter the number of retransmissions of each packet, called $N$. We investigate the impact of two values of $N$ on the reliability of Glossy, measured as the packet reception ratio (PRR). During one communication round, every node in the network initiates in turn a Glossy flood and all the other nodes log whether they successfully received the packet. This is repeated with for $N = \{1, 2\}$. In addition:

- The evaluation runs on TelosB motes[10] (26 nodes);

[6] casestudy_congestion-control.ipynb
[7] pantheon.stanford.edu/result/6539/
[8] github.com/StanfordSNR/pantheon
[9] casestudy_glossy.ipynb
[10] www.advanticsys.com/shop/mtmcm5000msp-p-14.html

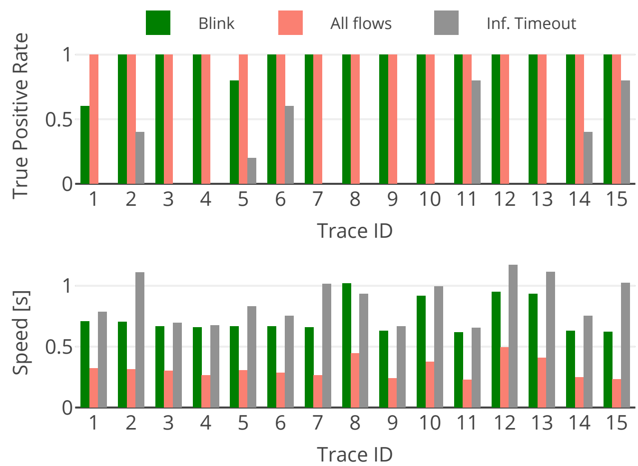


Fig. 8. KPIs for Blink's performance evaluation. *95% CI on the median. Internet trace IDs listed in [37].*

- The motes use radio frequency channel 22 (2.46 GHz, which largely overlaps with Wi-Fi traffic);
- The payload size is set to 64 bytes.

**Data collection.** We perform the experiments using the FlockLab testbed [46]. For both settings of the number of retransmissions $N$, we perform 24 randomly scheduled tests per day during 7 consecutive days. The data we collected are available on Zenodo [7].

### B.3 Failure Detection

**Reproducing the case study.** The entire case study is described in detail in a Jupyter notebook[11] which is available in the *TriScale* repository [5].

**Evaluation scenario.** This case study re-uses one of the evaluation scenarios from the original Blink paper (§ 6.1 in [37]). It considers 15 publicly available real Internet traces [22, 24]. For each trace, 30 prefixes are randomly selected among those that contain sufficiently many active flows. For each prefix, the characteristics of the traffic are extracted and used to run simulations where traffic sources generate flows exhibiting the same distribution of parameters than the one extracted from the real traces. Artificial failures are introduced in the simulation, which Blink tries to detect. Blink is compared against two baseline strategies:

- *All flows*, which monitors up to 10k flows for each prefix and reroutes if at least 32 of them sees retransmissions within the same time window. This strategy provides an upper-bound on Blink's ability to reroute upon actual failures but ignores memory constraints.
- *Infinite Timeout*, which is a variant of Blink where flows are only evicted when they terminate (with a FIN packet) and never because of the flow eviction timeout. This strategy tests the effectiveness of Blink's flow eviction policy.

**Data collection.** The authors of Blink kindly provided the data they collected for the original paper [37]. The data are now available on Zenodo [7].

[11]casestudy_failure-detection.ipynb

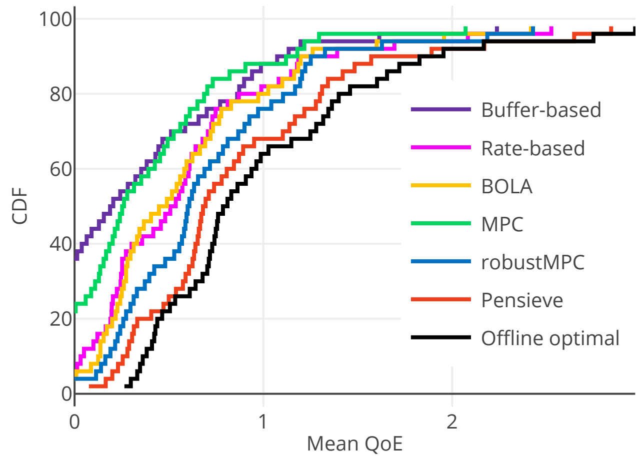


Fig. 9. 95% CI on the CDF of various adaptive bitrate algorithms.

**Evaluation objectives.** Each prefix is used to generate five failure scenarios, based on which we compute two metrics: (i) the true positive rate (TPR), *i.e.*, the ratio of failures that Blink successfully detects (out of 5); (ii) the median rerouting speed, *i.e.*, the time Blink takes to reroute traffic once it detects the failure. For both metrics, we use the 95% CI on the median as KPI, computed over the set of prefixes for each Internet trace.

**Results.** Blink achieves a TPR KPI of one for all the Internet traces, with a rerouting speed ranging between 0.5 to 1 s (Fig. 8). Hence, we can claim with 95% confidence that these are the minimal performance expected for Blink for any random set of prefixes within each of the Internet trace.

## B.4 Video Streaming

**Reproducing the case study.** The entire case study is described in detail in a Jupyter notebook[12] which is available in the *TriScale* repository [5].

**Evaluation scenario.** This case study re-uses one of the evaluation scenarios from the original Pensieve paper (§ 5.2 in [48]). Specifically, it compares Pensieve against pre-existing adaptive bitrate algorithms using different quality of experience (QoE) metrics. The comparison is performed using the MahiMahi [53] network emulator by replaying a set of synthetic traces generated from real-world broadband datasets. We consider the set of traces generated from the FCC dataset;[13] these traces were created by the Pensieve authors by concatenating randomly selected traces from the "Web browsing" category in the August 2016 collection. There are multiple definitions of QoE: we consider the "linear" one (see [48] for details).

**Data collection.** The authors of Pensieve were not able to provide the data they collected for the original paper [48]. Consequently, we retrieved the QoE data directly from the paper plots using a web-based application.[14] The data we retrieved are available on Zenodo [7].

[12] casestudy_video-streaming.ipynb
[13] Federal Communications Commission. https://www.fcc.gov/reports-research/reports/
[14] apps.automeris.io/wpd/

**Evaluation objectives.** From the QoE metric values, we compute the 95% CI (lower-bound) for the $\{2, 4, 6 \ldots 98\}$th percentiles, based on which we obtain a 95% CI for the entire CDF of QoE for the different algorithms.

**Results.** Fig. 9 shows the 95% CI CDFs computed for the linear QoE metric. The 95% CI are relatively close to the empirical CDFs, as illustrated in Fig. 7, which shows both the empirical CDF and its 95% CI for Pensieve (the same applies to all algorithms).