

A new methodology for representation of TCP performance in multiconnection environments

Chardalias Nikolaos

nihardal@ceid.upatras.gr

Georg-August Universität
Göttingen

Introduction (1)

- Evaluation of the different versions of the most common Internet transport protocol, TCP
- The most existing methodologies and representation tools are adapted to a simple scenario with few long-lived connections
- Thus, simulation of scenarios involving a large number of connections – closer to reality – is necessary

Introduction (2)

- In recent years great effort has been devoted to improve simulations models and scenarios in order to extract better and more general conclusions
- Different traffic models for Internet applications have been proposed and can be used in TCP performance tests
- This allows to perform comparisons for different TCP flavors and configuration parameters in more realistic scenarios

TCP Performance Graphical Representation (1)

- Most common approach: *time vs sequence number* plots, where any given long lived connection is traced and can be analysed in detail
 - Advantage: they show the per-packet behaviour of TCP and this allows the analysis of the TCP mechanisms performance to cope with losses and retransmissions and also allows the validation of these mechanisms
 - Disadvantage: representation of many connections in one plot is possible but the plot becomes blurred.

TCP Performance Graphical Representation (2)

- This talk focuses on a new graphical representation tool to evaluate the results of TCP simulations
- Simulations involve a realistic Internet model where many long-lived and short-lived connections coexist.
- All connection sizes can be suitably represented in a graph for correct interpretation of the effect of a given TCP version or flow control mechanism

Simulation Framework (1)

- The simulation scenario (fig. 1) consists of a set of web clients connected to the Internet through an access network
- These clients extract information from several servers during browsing
- The web traffic is a mix of short and long lived flows (also known as mice and elephants)
- Simplification: a constant Internet latency is assumed for all connections




Figure 1: Network Topology

Simulation Framework (2)

- A structural traffic model is used to achieve a WWW client-like traffic. The activity of a WWW user session is reproduced by means of a hierarchical multilayered structure. These levels are:
 - Session Level: session is the period of time that the user is surfing the Internet (birth-death simulation)
 - Page Level: the user must download any page he wants to visit (modelled with ON(request)/OFF(end))
 - Connection Level: TCP connections between the browser and the HTTP server in order to start the retrieval of the HTML file of the page (modelled with birth-death process)

Simulation Framework (3)

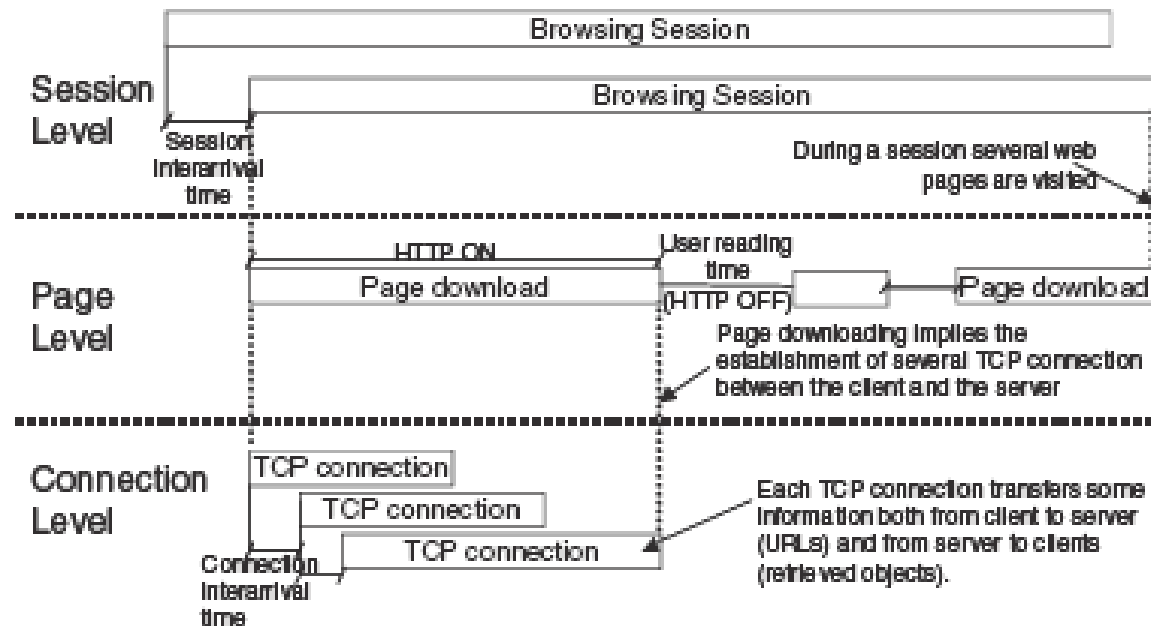


Figure 2: WWW traffic model structure

Goodput Based TCP Analysis (1)

- Evaluation criteria:
 - Throughput: the ratio between the total amount of bytes generated by a connection and the total amount of time. Information about the network resource usage per time unit.
 - Goodput: the ratio between the amount of application data transferred through a connection and the time necessary to complete their transmission.
 - Efficiency: the ratio between the total volume of application data transferred through a connection and total amount of bytes generated by the connection
 - Fairness: reflects how fairly bandwidth is shared among coexisting connections

Goodput Based TCP Analysis (2)

- In TCP simulations, the goodput of a connection depends on its size, because of the *slow-start* TCP mechanism
 - Small sized connection are heavily affected by the *slow-start* as they are usually finished before further control can be applied
 - Larger connections are not affected by that

Hence, correct evaluation must of a connection implies a two-component vector: *goodput* and *connection size*

Goodput Based TCP Analysis (3)

- A set of 2-dimensional vectors is usually represented with scatter plots (goodput in one axes, size on the other and the vector is a point in the R^2 plane).

The drawback is that the plot can not be easily readed when many vectors are graphically represented. To solve this problem, the idea is to use joint probability density function (joint pdf) of both magnitudes. That shows the probability of occurence of a given vector (g,s) where g is goodput und s the size. The joint pdf is extracted from the bidimensional histogram h_{GS} .

Goodput Based TCP Analysis (4)

- In order to reduce the computation costs, the histogram is computed for a goodput range defined by a minimum value G_{\min} and a maximum one G_{\max} and a size range given by S_{\min} and an S_{\max} .
- Most vectors on the h_{GS} are clustered on the bottom left part, as the mice are far more frequent than elephants in Internet. Thus, to further improve the visibility of the plot, the conditional probability density function of the goodput for any given connection size is used ($f_{G/S}(g,s)$). The conditional pdf is also estimated by means of a conditional histogram $h_{G/S}$.

Isopercentile Plots

- Present several goodput percentiles for each range of connection sizes, the conditional probability density function.
- The closer the curves are for a given connection size, the more connections achieving a similar goodput for that size range and the higher the conditional pdf is.
- Most isopercentile plots present a sigmoidal shape: small connections (affected by slow-start) have low goodput and goodput grows with the connection size until it starts to decrease in a size range.
- The slopes of the decreasing and increasing parts of the goodput give important information about the TCP version or configuration.

Experiment

- To show the usefulness of these plots, different TCP flavors are compared in a heterogenous environment where all versions coexist.
- 4 TCP flavors: New Reno (NR), New Reno with Limited Transmit (LT), FACK (F), FACK with Limited Transmit (FLT)
- Figures show 5%, 10%, 20%, 30%, 40%, 50%, 60%, 70%, 80%, 90%, 95% conditional goodput percentiles
- First scenario: queue of 20 KB and 1,3% packet loss.
- First scenario: queue of 10 KB and 4% packet loss.

Experiment – Results




Figure 3: Isopercentile plots for TCP flavor evaluation in 20 kB queue buffer scenario




Figure 4: Isopercentile plots for TCP flavor evaluation in 10 kB queue buffer scenario

Experiment – Analysis

- NR and LT show similar performance as seen in the plot. However, lower percentiles are higher for LT than NR for medium and large connections.
- Plots of F and FLT are bent to the right, meaning that *elephants* receive higher goodputs.
- The variability of goodputs is larger for versions F and FLT than for NR and LT.
- Largest connections in FLT present a lower variability than in F.

Evaluation of the Plots

- Even though quite realistic, this simulation is far from optimal.
- However, the representation of the TCP performance is given by a simple and intuitive way.
- Drawback: given the different goodput scale of *mice* and *elephants*, the behaviour of the smallest connections cannot be clearly appreciated. A different plot is required to observe the small connections and another one for the medium and large connections.

Radar Plots (1)

- Improve the isopercentile plots and also take into consideration the parameters of the configuration of the TCP flavors
- For comparison purposes (where many TCP flavors or configurations need to be compared), a transformation of the isopercentile plots is suggested.
- The goodput percentiles of each simulated TCP correspond to the radii of the figures, and different versions or configurations are given in different angles.

Radar Plots (2)

- The closer a line in a figure is to the external circumference, the better the goodput achieved for that connection size.
- This representation allows an easy and intuitive comparison analysis for each connection size.
- Percentiles are represented by a polygon: the more regular it is, the more similar the behavior of all evaluated TCPs.
- Convex vertexes correspond to worse versions and concave vertexes to the best ones.

Experiment

- Same experiments as before but more TCP versions are used: Tahoe with (T) and without fast-retransmit (TNF), the Improvement (H96), Reno (R), New Reno (NR), New Reno with Limited Transmit (LT), SACK (S), FACK (F) and FACK with Limited Transmit (FLT).
- As before two scenarios, with queues of 20 and 10 KB and packet loss of 1,3% and 4% respectively.
- The results are given for a few representative connection sizes.

Experiment – Results




Figure 5: TCP flavour performance comparison in 20 kB queue buffer scenario using radar plots




Figure 6: TCP flavour performance comparison in 10 kB queue buffer scenario using radar plots

Experiment – Analysis

- All TCP flavors behave similarly for small connections as they are only controlled by the slow-start stage.
- In larger connection, TNF and R are worse as expected, because they do not present efficient recovery mechanisms for multiple losses.
- For largest connections, best behavior have the SACK, F and FLT because of the selective recovery.
- Almost the same results with the second scenario too.
- Many losses lead to irregular plots (differences between TCP flavors can be seen easily under heavy losses).

Improvement plots

- Improve the main drawback of radar plots: when many percentiles are represented, it is not obvious to quantitatively appreciate the improvement of a given percentile with respect to the others.
- They present a given goodput percentile for different connection sizes in all the evaluated TCP versions.
- Better visualization of different connection sizes in the same plot but absolute performance information is lost, so they are meant to be used in combination with the previous ones.

Experiment – Results




Figure 7: Improvement plots for TCP flavor performance comparison in 20 kB queue size scenario




Figure 8: Improvement plots for TCP flavor performance comparison in 10 kB queue size scenario

Experiment – Analysis

- FLT and F have the best behavior for all represented percentiles.
- Interesting that for medium sizes LT behaves better than more complex TCP versions like F or S.
- LT and NR behave similarly in all cases.
- The worst versions are obviously R and TNF.
- In the scenario with queue of 10 KB where there are more losses, the limited transmit mechanism is very significant to improve less favored connections.

Conclusions (1)

- The idea is to evaluate TCP in terms of the goodput achieved by different connections in the access link, including *mice* and *elephants*.
- Three plots are suggested, according to the goal of the analysis: *isopercentile plots* to evaluate a single TCP version or configuration, *radar plots* to compare different versions or configurations and *improvement plots* to present the performance of different versions with respect to a baseline one.

Conclusions (2)

- The experiments have been performed in a framework designed to simulate a www scenario where different TCP versions coexist.
- Connections are affected by queue overflow losses, depending on the queue buffer size.
- The proposed graphical representation methodology has proven useful to represent the results and visually understand them.
- This methodology is not constrained to the analysis of simulation scenarios but can be applied to massive data collected in real scenarios

Further work

- Further work will focus on evaluating the impact of TCP configuration parameters and enhancements in more realistic simulation frameworks.

- Thank you for your attention.
- Time for discussion now!!

