On the relationship between fundamental measurements in TCP flows

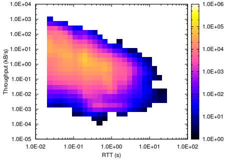

Richard G. Clegg (richard@richardclegg.org),João Taveira Araújo, Raul Landa, Eleni Mykoniati, David Griffin, Miguel Rio,
University College London, Department of Electronic Engineering

Talk to ICC Budapest 2013

(Prepared using LaTeX and beamer.)

Fundamental relationships within TCP flows

Problem Statement

Padhye et al – bandwidth (throughput) of TCP flow at equilbrium:

$$T = \frac{1}{D}\sqrt{\frac{3}{2bp}} + o(1/\sqrt{p}),$$

where D is RTT (delay), p is the probability of packet loss and b is a fixed TCP parameter.

- Result (simplified version presented) is from mathematical model with many assumptions.
- Subsequent work generalises and improves – basic inverse dependence on RTT and $\sqrt{p}$ remain fundamental.
- Is this true of real data?

Data and analysis approach

- Basic approach – use lots of freely available packet traces.
- Test both diverse data sets and similar data sets.
- Reconstruct TCP flows – calculate RTT, loss etc. Fit formulae relating these quantities.
- Data used CAIDA (US based data) MAWI (Japanese based data):
  - CAIDA OC48 Traces (2002) — 3 hours of data: 1.4 billion packets originally 876GB of data.
  - CAIDA OC192 (2011A) — 26 minutes of data: 1.3 billion packets originally 662GB of data.
  - CAIDA OC192 (2011B) — 14 minutes of data: 0.927 billion packets, 582 GB of data.
  - CAIDA OC192 (2012) — 29 minutes of data 1.6 billion packets and 1,120 GB of data.
  - MAWI (2006–2012) — 15 minute samples once per month, 1.36 billion packets and 982 GB of data.

Fundamental relationships within TCP flows

- In reality very little TCP is really TCP in the old-fashioned sense.
- TCP can be application limited (YouTube).
- TCP can be sender or receiver window limited.
- TCP can be limited by middleboxes.
- Ignoring all of this, what is the best relationship which ties network parameters to TCP performance.
- Step 1: graphically investigate the relationships in the data sets.
- Step 2: statistically fit equations which relate the parameters: throughput, loss, RTT, flow length.
- Consider subsets of data to ask questions about equilibrium and transient behaviour.

Data processing/filtering

- To get accurate RTT estimates only two-way data is considered.
- To get accurate RTT estimates only two-way data is considered.
- RTT can be inferred from SYN/SYNACK/ACK handshake.
- RTT can also be inferred from data transfer when data in both directions.
- OC48 and MAWI both directions seen majority of time. OC192 less so.
- Truncation effects mitigated by removing flows do not seem to end within lifetime of capture file.
- Starting point is to visualise correlations in data.
- Most interesting visualistaion comes from 3d histograms.

Visualising correlations throughput/RTT

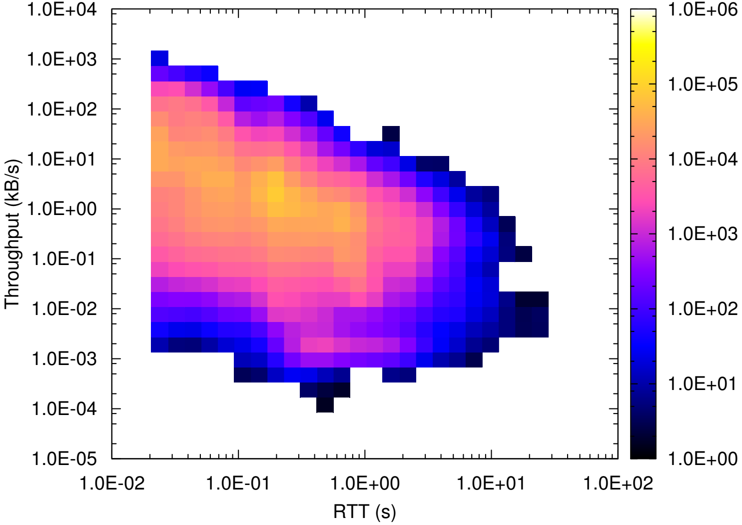

OC48 — relationship between throughput and RTT

Visualising correlations throughput/loss

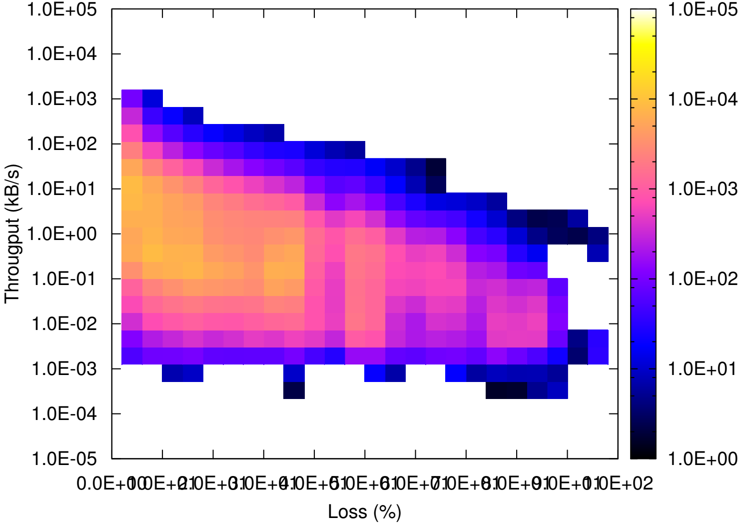

MAWI — relationship between throughput and loss

Visualising correlations – throughput/packets

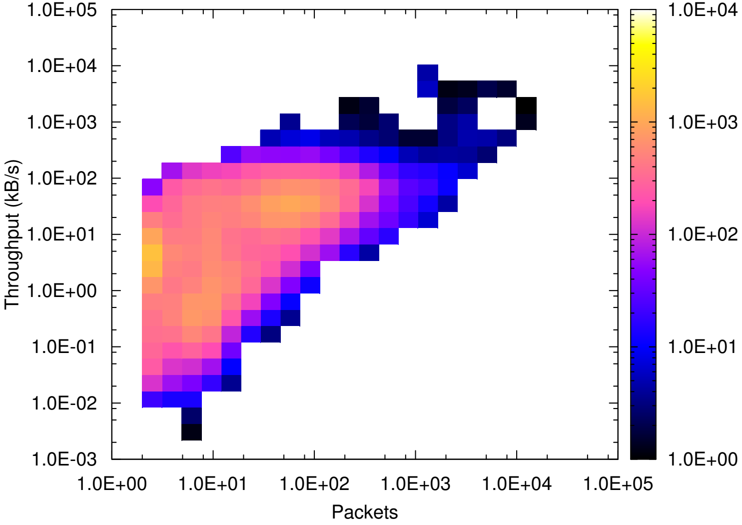

OC192 2012 — relationship between throughput and number of packets in flow

Fitting a Linear Model

- Variable Y is observed variable to be explained in terms of variables X_1, X_2 etc.
- Assume a linear relationship $Y = X_1 + X_2 + \cdots + \varepsilon$ where $\varepsilon\ N(0, \mu)$.
- Want to find β parameters to minimise the error term.
- Fit log of data and use exponential transform to get $T = \beta_0 D^{\beta_1} p^{\beta_2} \varepsilon'$ where ε' is mean 1, lognormal).
- With $\beta_1 = -1$ and $\beta_2 = -0.5$ this is $T = \beta_0 / D\sqrt{p}$ (and error term).
- Goodness of fit judged by R^2 value where $R^2 = 1$ is perfect and $R^2 = 0$ is no fit at all (amount of variance "explained" by model).
- Taking logarithms a problem for loss as sometimes $p = 0$ – use instead $\log p + p_m$ where p_m is a fitted offset parameter.

CAIDA OC192 2012 data

Model for T	R^2	Note
$15.7D^{-0.94}(p+p_m)^{-0.563}P^{0.456}$	0.641	$p_m = 0.105$
$77.2D^{-0.975}P^{0.455}$	0.635	
$316/(D\sqrt{p+p_m})$	0.0227	$p_m = 0.105$

- Excellent fit to data.
- Loss p slightly improves model but not much.
- Best model is approx $T = k\sqrt{P}/D$ where k is constant.

CAIDA OC48 data

Model for T	R^2	Note
$102D^{-0.929}(p+p_m)^{0.391}P^{0.339}$	0.362	$p_m = 0.105$
$29.7D^{-0.89}P^{0.354}$	0.35	
$193/(D\sqrt{p+p_m})$	0.207	$p_m = 0.105$

- Weaker fit to data but not bad for a simple model.
- Again p (loss) has little explanatory power.

CAIDA OC192 2011A data

Model for T	R^2	Note
$0.712D^{-0.665}(p+p_m)^{-0.661}P^{0.429}$	0.454	$p_m = 0.105$
$4.62D^{-0.698}P^{0.41}$	0.448	
$251/(D\sqrt{p+p_m})$	0.109	$p_m = 0.105$

- Reasonable fit to data.
- Again loss p not much help.
- Best model approx $T = kP^{0.4}/D^0.7$.

CAIDA OC192 2011B data

Model for T	R^2	Note
$21.5D^{-0.924}(p+p_m)^{-0.581}P^{0.419}$	0.616	$p_m = 0.105$
$156D^{-0.981}P^{0.386}$	0.611	
$562/(D\sqrt{(p+p_m)})$	0.19	$p_m = 0.105$

- Much better fit than 2011A.
- Best model approx $T = kP^{0.4}/D$.

MAWI data

Model for T	R^2	Note
$0.15D^{-0.664}(p+p_m)^{-0.416}P^{0.635}$	0.282	$p_m = 0.0132$
$0.648D^{-0.583}P^{0.576}$	0.332	$P > 1000$
$111/(D\sqrt{p+p_m})$	0.0904	$p_m = 0.105$

- Fairly weak fit to data.
- Perhaps because data over long time period.
- Best fit is only for long flows (more than 1000 packets).

Parameter dynamism

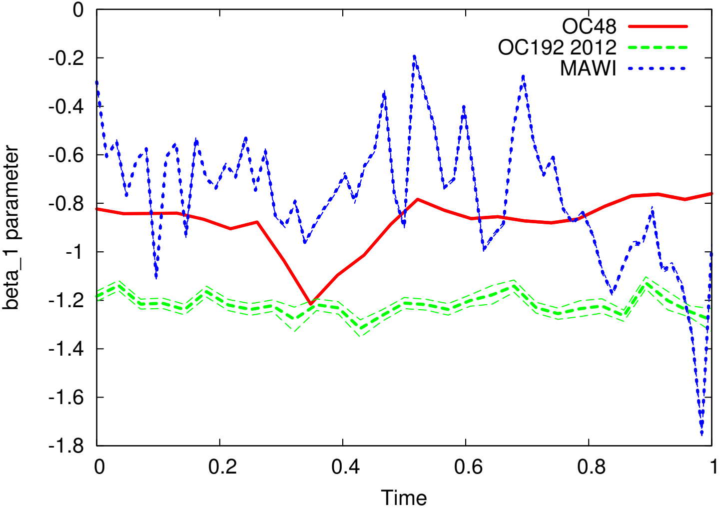

Evolution of β_1 parameter in model $T = \beta_0 D^{\beta_1}$ across normalised time

Conclusions and further work

- This work is just a starting point but appears to be the first to fit this type of model.
- Further work at UCL shows that for more than 50% of TCP flows in recent data controlling factors are not "standard TCP".
- However, these extremely simple models are often an excellent fit to data.
- In short traces the parameters remain surprisingly constant.
- Roughly speaking there is a $1/RTT$ relationship to throughput.
- The correlation with loss was very low.
- Length of the flow in packets was important (this has been observed by other researchers).

Questions?

?