Mobile Content Delivery Optimization based on Throughput Guidance

Péter Szilágyi (Nokia Networks)
Andreas Terzis (Google)
IETF 93, IRTF ICCRG, Prague, 22nd July, 2015

Disclaimer

- Early stages of understanding the benefits of exposing network-level capacity information at the end-to-end transport layer.
  - Information may cover part of the end-to-end path (e.g., radio downlink)
- Given the right constraints, we show the benefits of explicit capacity information
  - Radio link is the bottleneck
  - Caches deployed close to the users
  - One/Few flows per UE
- Danger of elevating congestion and packet loss in the general case
- Substantial work is necessary to evaluate when and how to best utilize available information
  - Safety mechanisms (e.g., generate signal only to whitelisted IP destinations)
  - Fallback schemes
  - Extending scope of capacity information (e.g., radio link vs. end-to-end path)
  - Nature of capacity information (increase/decrease vs. explicit rate, decrease only)

Overview

Throughput Guidance (TG): network-side information that enables data transfer and/or content optimization to:

(1) improve the efficiency of data transmission through mobile networks.

(2) improve the end-user experience.

Mechanism: expose the amount of bandwidth available at the (radio access) network to the data sources.

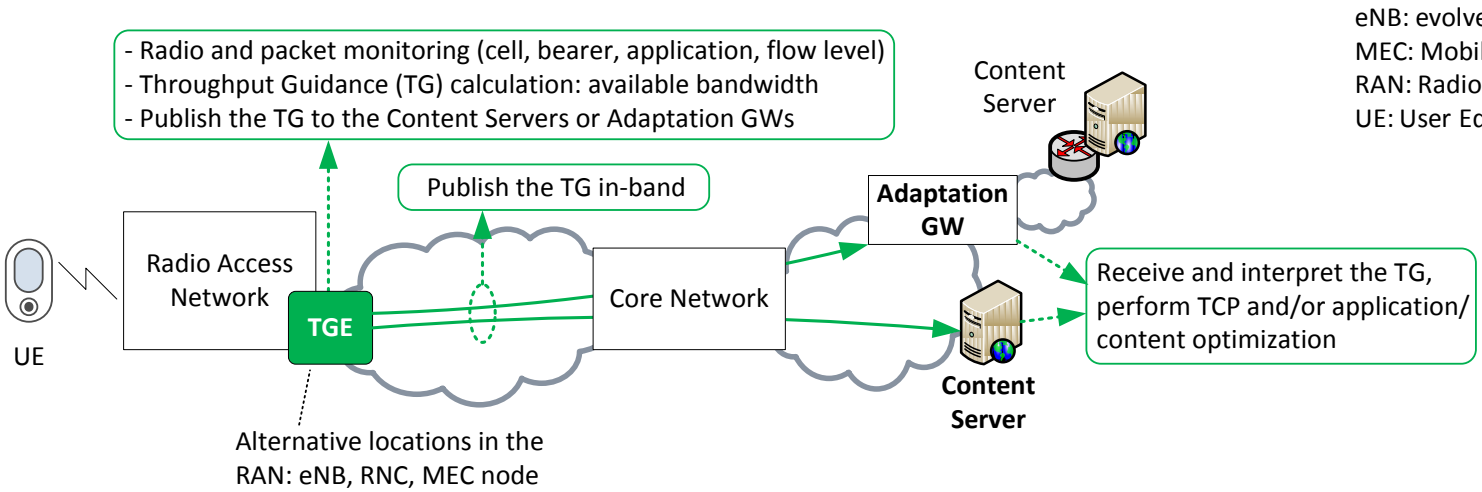

eNB: evolved Node B (base station)
MEC: Mobile Edge Computing
RAN: Radio Access Network
UE: User Equipment

Agenda

- Problem Statement
- Throughput Guidance Calculation
- Optimization based on Throughput Guidance
  - TCP optimization: slow start, congestion avoidance, packet pacing
  - Media (application layer) optimization: initial content selection, media rate adaptation
- Considerations
  - multiple flows within a bearer
  - co-existence of guided and unguided flows
  - reverting to standard TCP behavior
- Performance Evaluation: Field Trial

Problem Statement

1. TCP's own network probing mechanisms (a.k.a. congestion control):

   slow start, congestion avoidance, under-utilization or over-shooting of the available network resources

2. Content (application) layer:

   the bandwidth demand (media rate) of the selected content should not exceed the available resources to avoid customer experience degradations

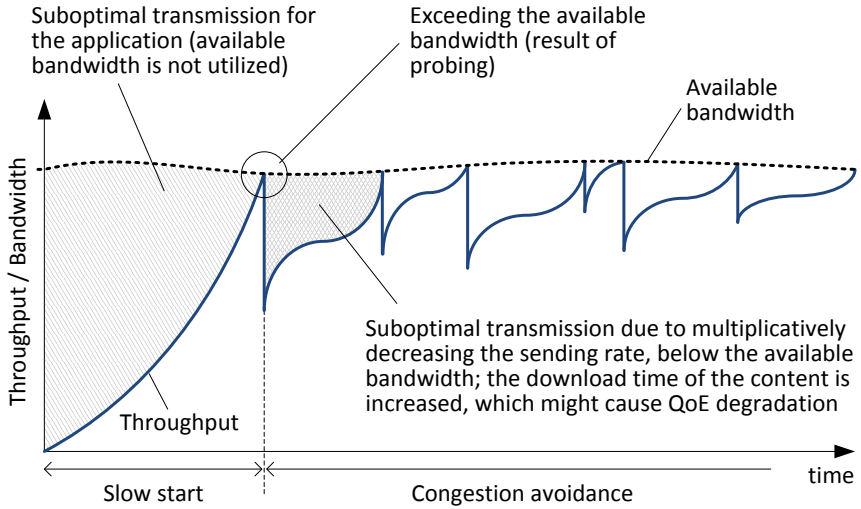

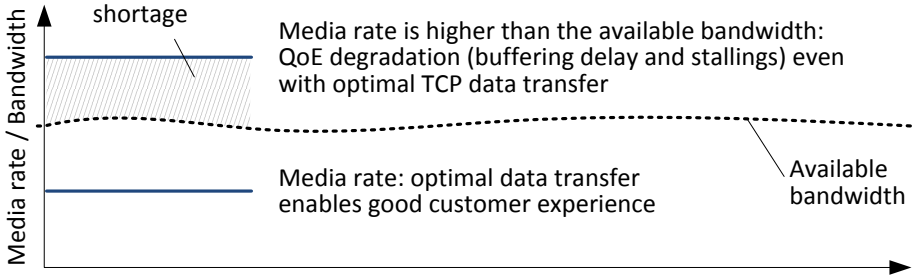

Note: the time scale and bandwidth on the figures are illustrative only.

Throughput Guidance Calculation: Per-Bearer Bandwidth

The TG is transmitted as long as the radio interface is the bottleneck in the network (detected by the TG Entity).

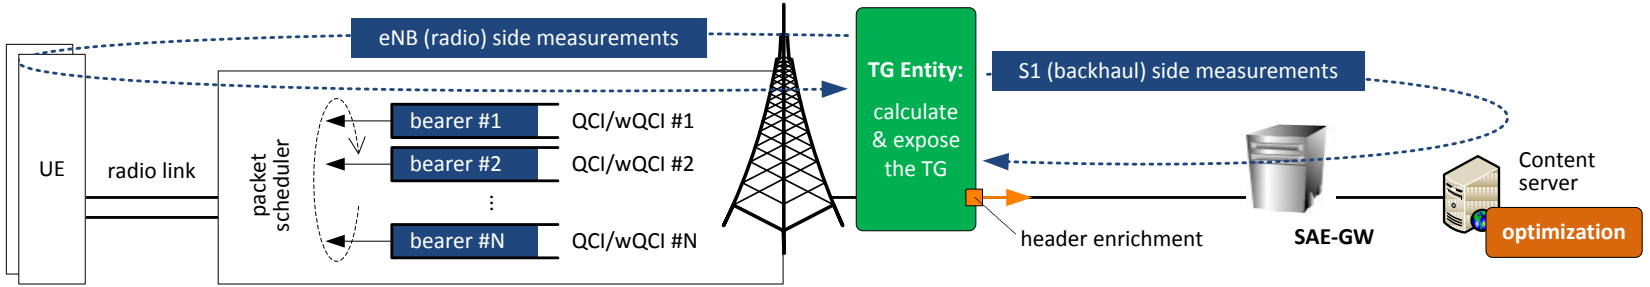

1 Detect radio/eNB side congestion

Radio bottleneck is detected by end-to-end, eNB side and S1 side throughput, loss, RTT and delay pattern monitoring.

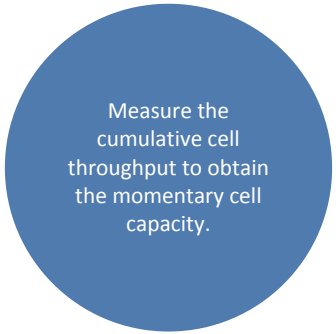

2 Calculate per bearer eligible throughput

Per-bearer fair share: radio channel and QoS aware segmentation of the momentary cell capacity.

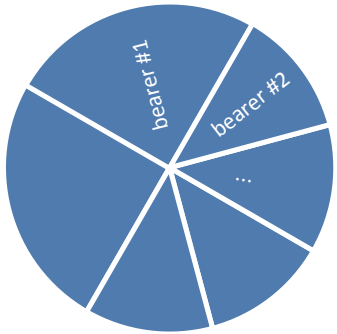

3 Compare the measured per-bearer throughput with the eligible throughput

Case #1:

The achieved throughput is fully utilized by the sources.

The TG equals the eligible throughput as the other bearers may claim their share anytime.

Case #2:

The eligible throughput cannot be achieved (e.g., due to poor individual channel quality) despite pending/ongoing data transfer (detected by the TG Entity).

The TG equals the measured throughput.

TCP Optimization: Overview

Scope: efficient utilization of the available network resources without having to probe for bandwidth and create congestion.

Method: the TCP sender adjusts its transmission rate to the value indicated by the TG.

Possible mechanism: set the congestion window (cwnd) as a function of the smoothed RTT (sRTT) and the latest TG:

$$cwnd = sRTT \cdot TG$$

Packet pacing: a technique used to mitigate the bursty transmission pattern of the TCP sources; with TG, the pacing rate can be adjusted directly.

Note: the advertised window should be taken into account (upper limit on the cwnd).

TCP Optimization: Slow Start (measurement example)

Start transmitting the data at the rate indicated by the TG value.

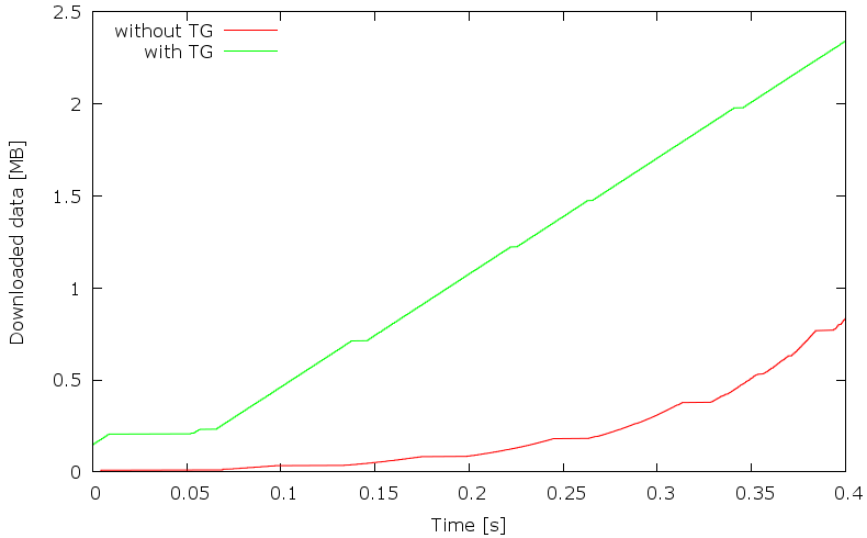

50 ms end-to-end RTT

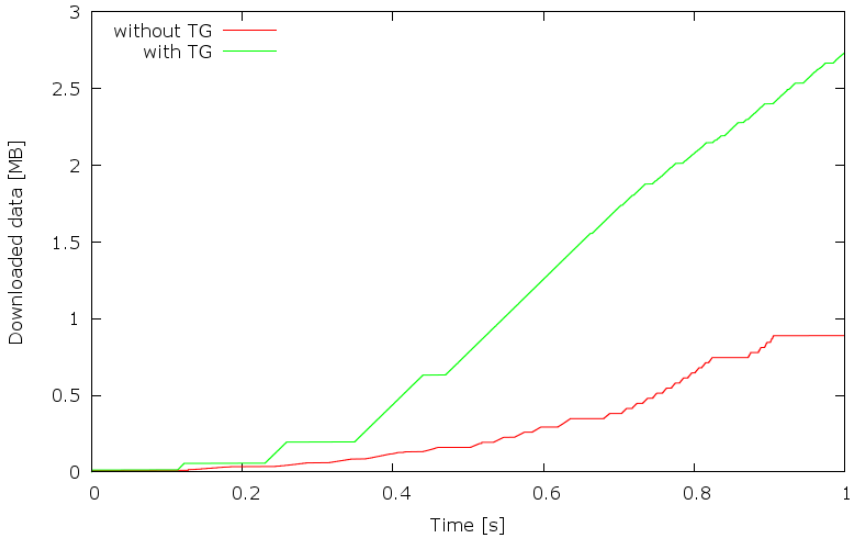

100 ms end-to-end RTT

Note: test results taken from LTE lab measurements; TCP: Linux Cubic (IW = 10); available radio bandwidth: 50 Mb/s

TCP Optimization: Congestion Avoidance – without TG (measurement example)

Regular TCP interprets all lost segments as a sign of congestion and multiplicatively reduces its transmission rate.

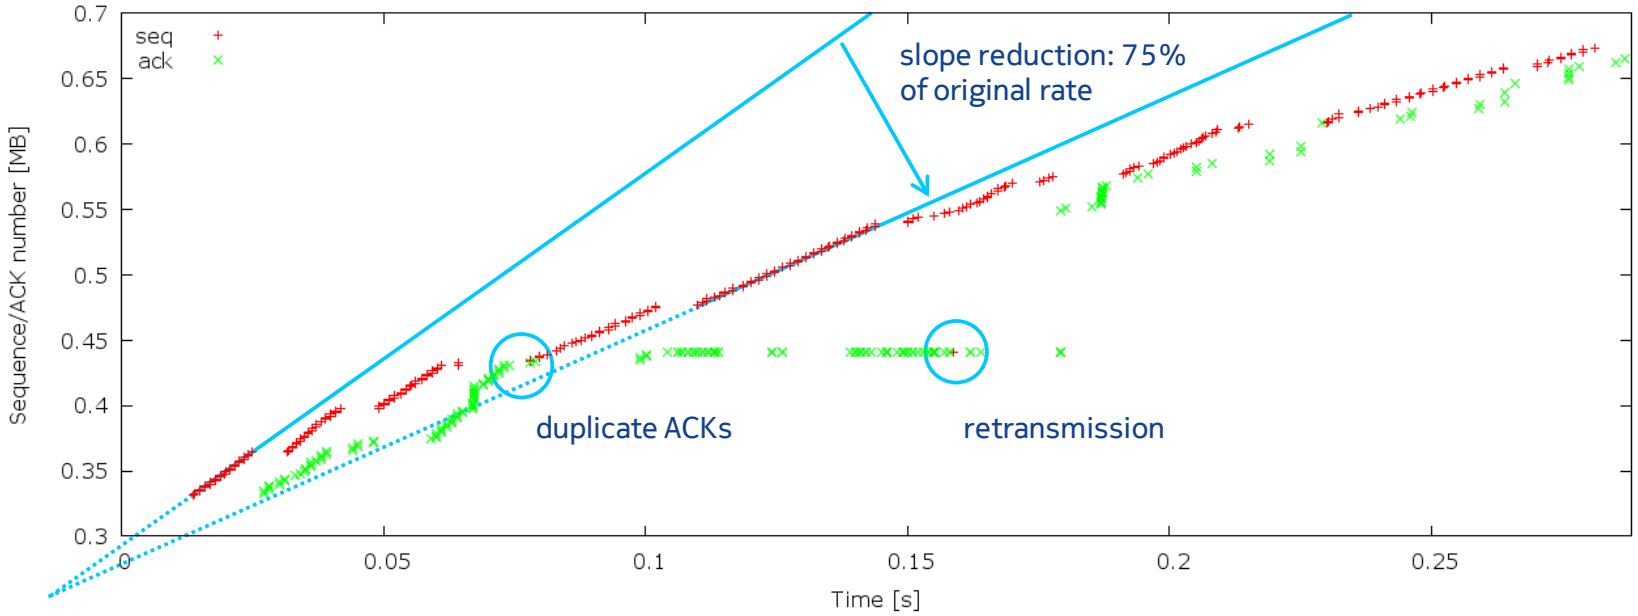

Regular TCP: fast retransmission and reduced rate according to TCP Cubic behavior

TCP Optimization: Congestion Avoidance – with TG (measurement example)

The TG Entity can indicate to the data source if a non-congestive loss happens on the radio; in that case, the TCP source may keep its transmission rate.

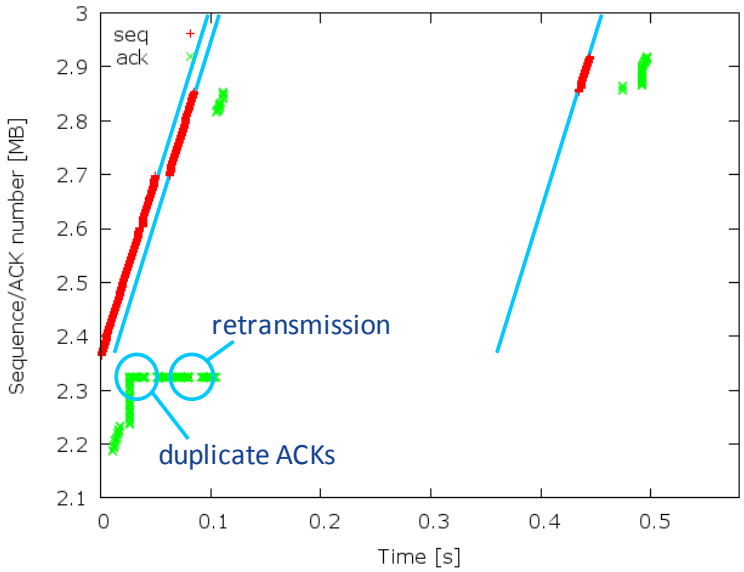

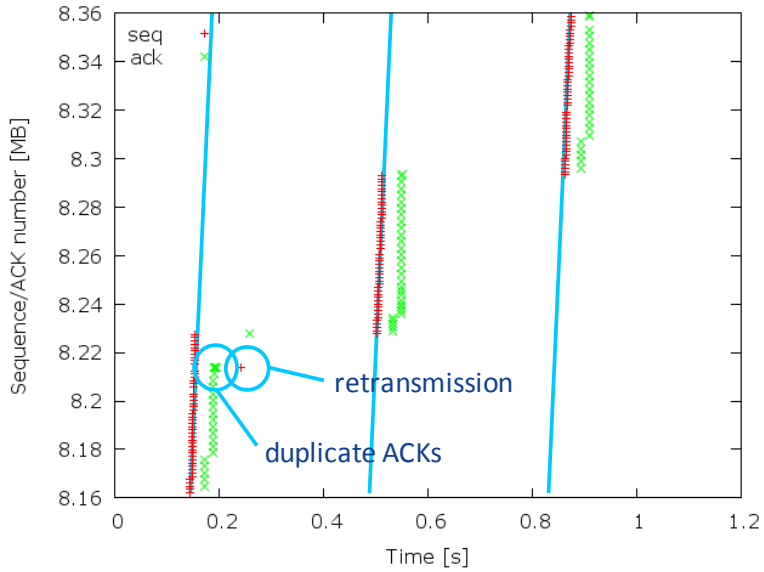

The transmission rate is kept at the value indicated by the TG.

Media (Application Layer) Optimization (simulation result)

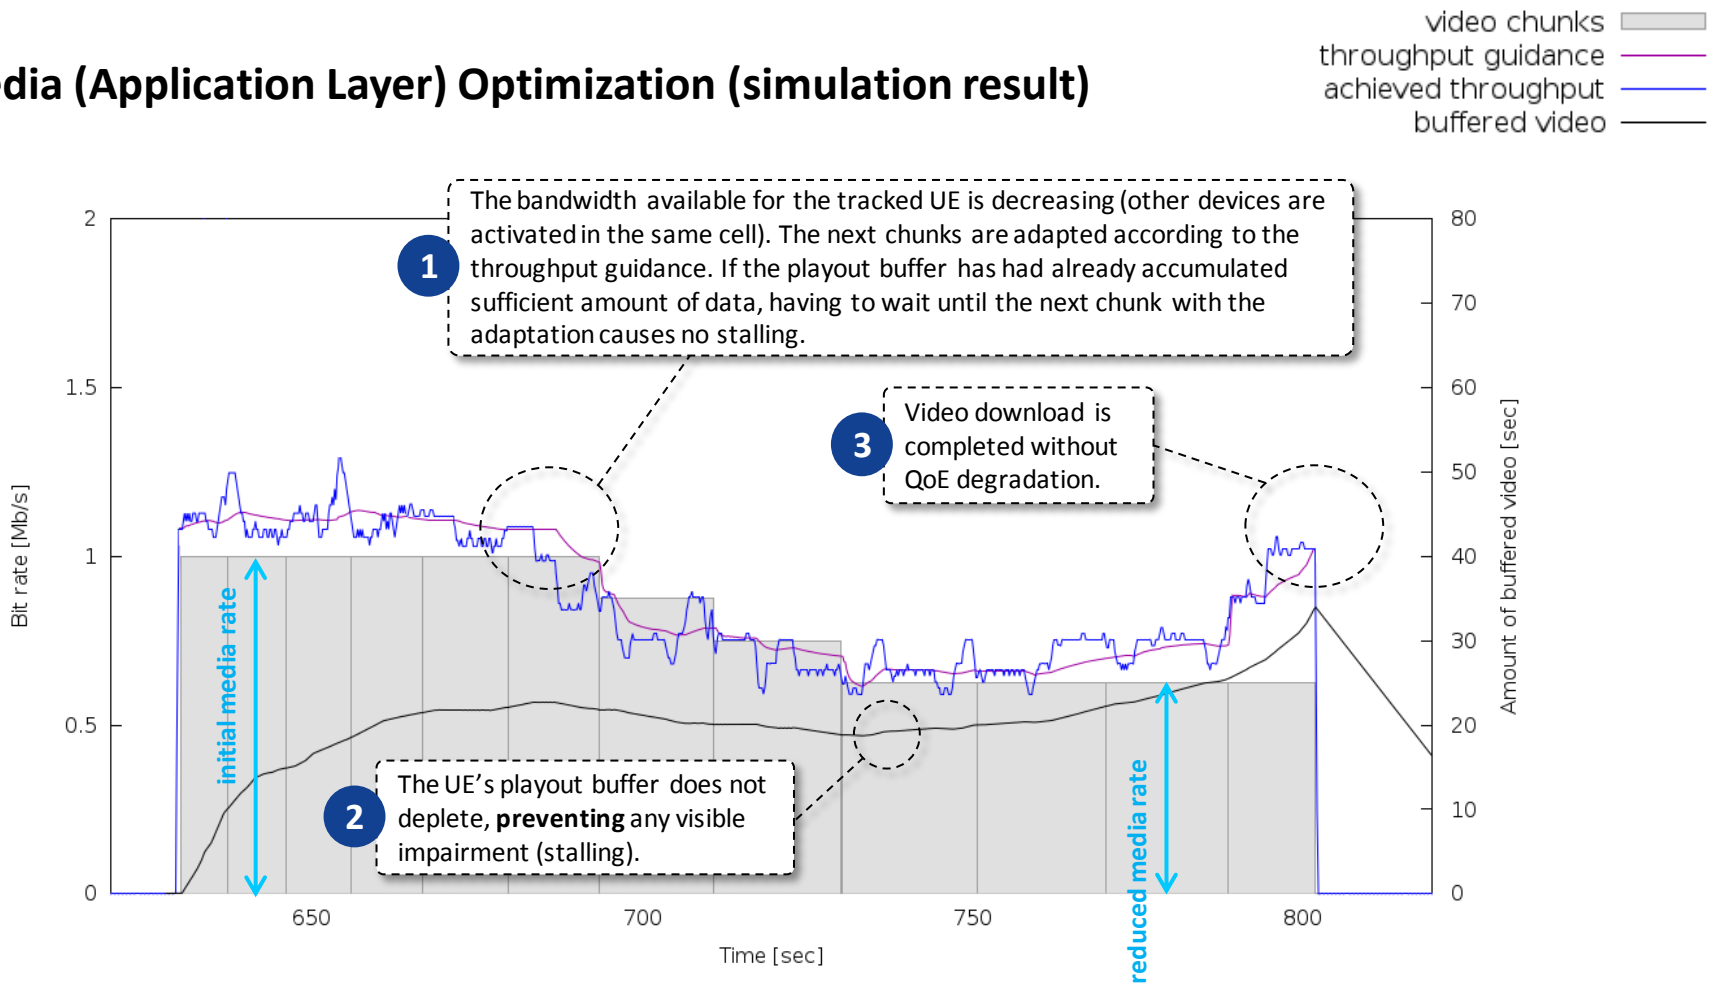

Note: Google is not using TG for application layer optimization.

Considerations for Throughput Guidance Calculation

1. What if there are multiple guided flows within the same bearer?

2. What if there are both guided and unguided flows within the same bearer?

3. What if the radio interface is not the bottleneck?

Consideration: Multiple Guided Flows within the Bearer

The per-bearer bandwidth is usable for optimization in case

(1) there is a single active flow in the bearer

(2) there is a single optimization entity terminating (or managing) all flows within the bearer

Otherwise, the TG Entity has to calculate a per-flow TG by dividing the per-bearer available bandwidth among the active flows and sending each TCP source its own fair share as the TG.

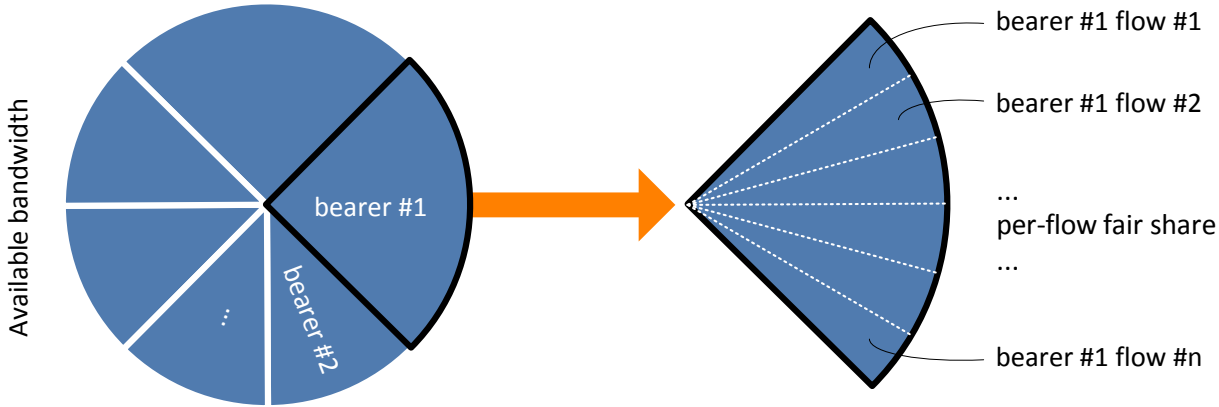

Consideration: Co-existence of Guided and Unguided Flows within the Bearer

Separate the optimized (guided) and non-optimized (unguided flows) and enforce the eligible share of each aggregate.

(1) protect unguided flows from the guided ones

(2) ensure that the bandwidth advertised via TG is available (not impacted by the TCP probing and overshooting of other flows)

Possible mechanism: intra-bearer WFQ scheduler, dynamically adjusted weights

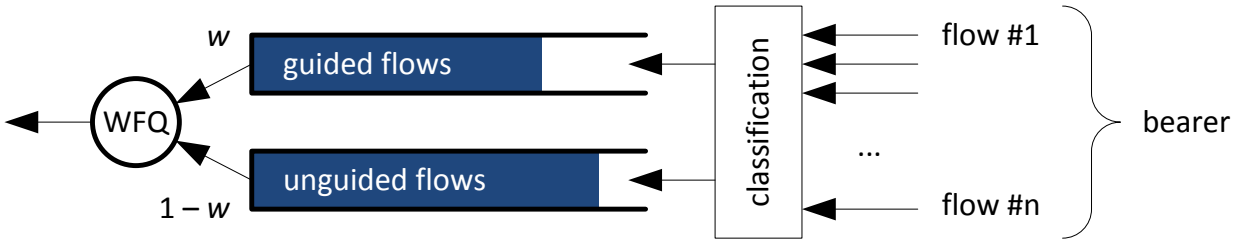

Consideration: Reverting to Standard TCP Behavior

The TG value may not be available at the TCP source for various reasons:

(1) TG exposure failure (e.g., option header enrichment not possible)

(2) no valid capacity measurement available (e.g., when the system is deployed and initialized, has been idle for significant amount of time, or the bottleneck is not at the resource monitored by the TG Entity)

In case (2), the TG Entity may transmit an indication of the reason for the missing TG value.

The TCP source should revert to the standard TCP behavior until a new TG value is received.

Performance Evaluation: Field Trial

Live production LTE network. Google server placed close to packet core.

Collected parameters:

- **Join time** (also known as time to play): measured from the user selecting a video until the video starts playing.
- **#Formats:** the number of different formats (resolutions) used during a video playback.
- **Avg. Video Res.:** the time-weighted average of all video resolutions used during the playback.
- **Client BW:** the average download rate of the video session as reported by the client.
- **Re-buffer Time:** The total amount of time that the video re-buffers.

Performance Indicator	Difference of Averages (%)	Difference of 99th percentiles (%)
Join time	-8.0%	-12.4%
#Formats	+4.1%	+29.9%
Avg. Video Res.	+6.2%	+5.6%
Client BW	+0.7%	+8.0%
Re-buffer Time	-19.7%	-5.1%

Related Drafts

Mobile Throughput Guidance Inband Signaling Protocol
draft-flinck-mobile-throughput-guidance-02.txt

Requirements and reference architecture for Mobile Throughput Guidance Exposure
draft-sprecher-mobile-tg-exposure-req-arch-01.txt