B. Han et al, “An Anatomy of Mobile Web Performance over Multipath TCP” in CoNEXT15

Taeju Park, Buting Ma

Introduction

- Mobile devices are equipped with multiple network interfaces (WiFi, cellular)
- Many studies have discovered impact of the multi-path networking over MPTCP on the mobiles.
  - Performance
  - Energy consumption
  - Mobility
- How multi-path networking over MPTCP impacts on web browsing performance (especially in page load time (PLT)) on the mobiles?
  - Loading web page is a interleaving procedure between communication and computation.
  - Employs many short-lived TCP connections with only a few round-trips
  - Depends on the web protocol (e.g, HTTP, SPDY)

Questions

- How much PLT reduction is achieved by MPTCP, comparing to SPTCP?
- How the PLT reduction is changed under various network condition?
  - Latency
  - Loss rate
- How the web protocols impact on PLT with MPTCP/SPTCP?
  - HTTP/1.1
    - One request per one TCP connection
    - Uses a number of concurrent short-lived TCP connection
  - SPDY
    - Many request per one TCP connection
    - Multiplexes web objects into a single connection with a longer duration

Performance Measurement Tool

- Develop “tcpdump-mpw”
  - Extracts HTTP/SPDY request/response data from raw packet traces.
  - Input
    - ‘Pcap’ capture, SPDY proxy’s TLS private key
  - Process
    1. MPTCP subflow assembling
    2. MPTCP logical connection assembling
    3. TLS/SSL decryption
    4. HTTP/SPDY parsing
    5. Web object information extraction
  - Output
    - A table providing details of HTTP/SPDY transaction (URL, raw size, decoded size, content type, expiration time, begin/end time and the amount of raw data transferred over MPTCP subflow.)

Performance Measurement Testbed

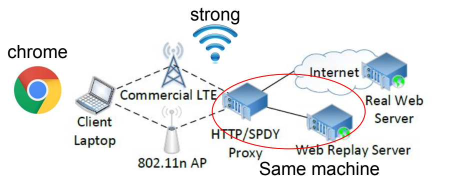

Figure 1: Our multipath measurement testbed.

- HTTP/SPDY Proxy
  - Today's website does not support SPDY
  - HTTP and SPDY proxies are co-located. Thus, we can compare two protocols fairly.
- DummyNet
  - Add delay, loss, or throttle throughput to the two wireless paths.

Automated Page Loading Experiments

- Measure 25 representative websites
  - Snapshot in August 2014 using replay server.
- Measure PLT
  - Time span between issuing the "landing page" request and reception of the "load event"
- Configurations
  - Cold-cache loading (caches are cleared)
  - High-power state (RRC_CONNECTED) for LTE interface
    - Avoid radio state transition delay
  - 100 successful measurement for each website
  - Use full version of website instead of mobile version

Website	# Obj	Size (KB)	# Txt Obj	# JS/ CSS	# Img Other	RTT (ms)
News	188.3	4450.8	33.8	67.7	86.8	20.9
Univ.	68.0	2583.9	3.0	16.0	49.0	30.1
AD	19.0	1459.6	3.0	6.0	10.0	62.0
News	201.2	3821.2	38.9	58.0	104.3	10.7
Tech.	67.0	2152.2	7.0	30.0	30.0	22.2
Video	93.6	2662.9	11.0	32.6	50.0	71.9
Football	99.1	2456.2	9.0	39.1	51.0	20.9
SHOPN	52.2	1000.6	5.0	12.0	35.2	8.97
Radio	66.2	2453.0	16.0	23.0	27.2	3.65
Wiki	28.0	601.2	1.0	6.0	21.0	8.70
Search	2.0	59.9	1.0	0.0	1.0	8.55
ENT	60.9	2170.6	13.4	25.1	22.4	71.8
Image	91.0	3275.2	3.0	15.0	73.0	6.20
Sport	119.0	2651.4	13.0	26.0	80.0	73.4
Social	69.0	1700.2	3.0	17.0	49.0	9.48
Movie	39.0	845.7	6.0	3.0	30.0	9.09
Weather	143.9	3814.2	21.4	39.7	82.7	11.7
RE	26.0	894.2	2.0	10.0	14.0	69.9
Airline	68.0	2148.6	4.0	23.0	41.0	3.15
SHOPN	48.0	2288.9	1.0	24.0	23.0	14.2
Gov.	68.0	2774.7	4.0	13.0	51.0	8.20
Travel	21.0	2000.4	1.0	2.0	18.0	3.85
Dict.	72.4	2223.1	10.8	31.7	29.9	6.30
Finance	39.7	1988.1	2.6	15.1	22.0	6.17
Market	49.0	2032.8	2.0	16.0	31.0	2.97

Measurement

Path	Downlink (Kbps)	Uplink (Kbps)	RTT (ms)
WiFi	7040	2020	50
LTE	9185	2286	70

Table 2: Throughput and RTT for baseline experiments.

Apply Dummynet to artifact bandwidth and latency

Six tests:

SPTCP:

HTTP WIFI

SPDY LTE

MPTCP: (WIFI primary subflow)

HTTP

SPDY

Base-line

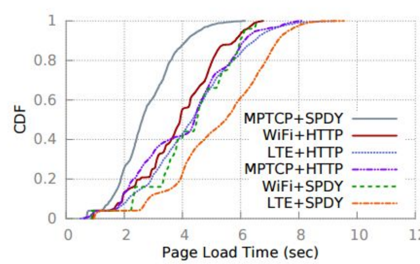

Figure 2: PLT distributions for baseline experiments (SPTCP and MPTCP).

Single-path:

HTTP better than SPDY

(Losses and large objects deteriorate SPDY)

Multi-path:

SPDY is improved greatly, but HTTP is not (even slightly worse on chrome)

Why HTTP doesn't benefit

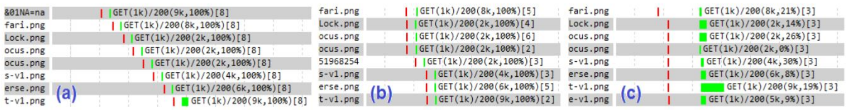

(a) HTTP over MPTCP (b) HTTP over SPTCP (WiFi) (c) SPDY over MPTCP.

Sequential behavior:

the browser gives that connection (i.e., Connection 8 in Figure 5(a)) a high priority based on various heuristics,
e.g. the browser estimates the connection's congestion window to be large due to its use of MPTCP, and thus the browser may think using that connection to transfer objects sequentially takes shorter time than using other idle connections

Larger cwnd ---> prefer one connection over others (HTTP unknow smtp)

Why SPDY benefits

First, SPDY suffers from bandwidth under-utilization (slow start in single connection), alleviated by MPTCP.

In contrast, most HTTP flows are short-lived, giving them fewer chances to use both paths: 100% SPDY use both WiFi and LTE, but only 47.4% HTTP use both

Second, SPDY's vulnerability to losses, which is again attributed to its usage of a single connection, is also mitigated by MPTCP.

If a loss happens on one path, MPTCP can use another one to minimize the impact of loss.

Packet Losses (SPTCP)

SPDY is worse than HTTP

In SPTCP, SPDY's single TCP connection reduces its congestion window, too conservative for non-congestion (random) losses that commonly occur in wireless networks.

In contrast, in HTTP/1.1, a non-congestion loss only affects one of its parallel connections, leading to a smaller impact.

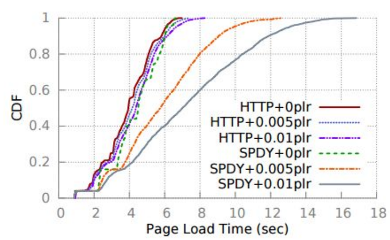

Figure 3: PLT distributions when additional losses added to WiFi (SPTCP).

Packet losses (MPTCP)

MPTCP prefers less lossy LTE path when losses occur over WiFi.

under 0%, 0.5%, and 1% additional PLR,
For SPDY, 62.2%, 66.3%, and 69.2% (median) of TCP payload is delivered over LTE.
For HTTP, 35.5%, 36.4%, and 37.0% on LTE, respectively,

because HTTP flows are very short and thus the LTE path has much smaller chances of being used (no time to adjust).

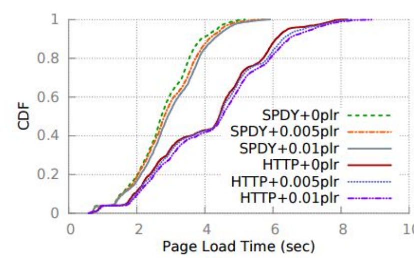

Figure 4: PLT distributions when additional losses added to WiFi (MPTCP).

Latency

over MPTCP, SPDY outperforms HTTP, even when additional latency is present.

For both HTTP and SPDY, delay on WiFi incurs more performance degradation than on LTE.

MPTCP prefer path with shorter RTT: 50+50/70; 50/70+50

For HTTP, originally, 62.3% (median) of downlink payload is delivered over WiFi. After adding a 50ms delay to LTE, the fraction remains similar at 69.0%, while adding a 50ms delay to WiFi dramatically reduces it to 37.1%

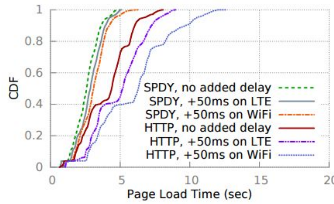

Figure 6: PLT distributions when additional latency added to each path (MPTCP).

HTTP over MPTCP is more sensitive to delays compared to SPDY over MPTCP.

network time of HTTP counts heavier, hence a proportional increase of network time impacts more

Congestion Control

- LIA (Linked Increase congestion control Algorithm)
- OLIA (Opportunistic LIA)
- Conventional CC (CUBIC) to each MPTCP subflow (more aggressive)

Small impact, smaller on HTTP

HTTP uses many concurrent short-lived connections (0.371s), most of which terminate at the slow start phase so CC does not yet have a chance to operate

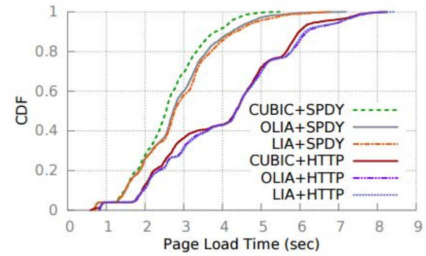

Figure 7: PLT distributions under different congestion control algorithms (MPTCP).

Real Websites

Proxy server go fetching instead of replay

1. Longer network time; 2. longer processing time (overloaded server); 3. several non-origin sources

disparities among the six schemes become smaller.

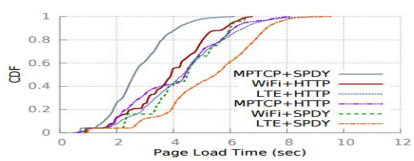

Figure 2: PLT distributions for baseline experiments (SPTCP and MPTCP).

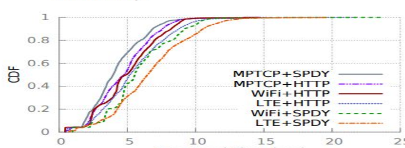

Figure 8: PLT distributions for real websites (SPTCP and MPTCP).

Summary and Recommendations

1. Multiplexing-based web protocols, such as SPDY and HTTP/2, gain more benefits from MPTCP than HTTP/1.1 does.

2. We found for Chrome, HTTP over MPTCP is sometimes worse than HTTP over SPTCP, possibly because MPTCP's interference leads to suboptimal TCP connection management decisions.

3. Since web browsing is usually delay-sensitive, when using the default scheduler of MPTCP, the PLT is largely determined by the path with a lower latency when both paths' bandwidth and loss do not differ too much.

4. Although MPTCP may significantly boost the performance of network transfer, when page loading is throttled by local/remote computation, the gain of enabling MPTCP may be limited.

Discussions

- Client Device
  - Use laptop with the full version webpages
  - Handheld device with the mobile version webpages may reduce local computation time
- Energy Consumption
  - Previous study shows that using LTE and WiFi simultaneously for file transfer increases overall energy consumption.
  - Does web browsing shows similar trends to file transfer?
    - Vs. Reduce PLT (reduce communication and computation time)
- Others
  - Combination of path's characteristic (loss, delay and congestion control)
  - Other browsers (Firefox, Safari and etc.)

More Discussions

- Does the HTTP/2 really shows same PLT distributions of SPDY?
  - Similar multiplexing mechanism
  - But, there are different things (e.g, header compression)
- How about 3G cellular (CDMA, GSM) instead of 4G cellular (LTE)?
  - Different speed
  - Different energy consumption
- They only provide average PLT distribution over all types of web pages.
  - Is the PLT distribution same for the different types of web pages?
    - Text-intensive
    - Image-intensive