WarpTCP - Solutions

WHITE PAPER

TCP De-Bottleneck - Challenges for Users, Enterprises & Providers

BADU networks

- Improving the way the world connects -

Sluggish Network - What are the Causes?

When facing a sluggish network, it is important to diagnose the root cause of the problem.

Vertically, the problem can be diagnosed according to the OSI model — among the layers, the Application (L-7), Transport (L-4), Media Access (L-2), and Physical (L-1) layers are the more important ones. The end-to-end throughput will be the minimum of the available throughputs in traversing these layers.

Horizontally, the problem can be diagnosed according to the devices and links in the end-to-end path: from one end host, through the various middle boxes and links, to the other end host. Again, the end-to-end throughput will be the minimum of the available throughputs in traversing these devices and links.

The real bottleneck in an application will be the same one, either from the horizontal or vertical analysis. While this is common knowledge among experts, often the real bottleneck is not carefully diagnosed. The issue is that all components in either the vertical or horizontal analysis interact with some components in a nonlinear fashion. It could happen that upgrading a non-bottleneck actually improves the end-to-end speed, thereby causing the non-bottleneck to be misidentified as the real bottleneck— when this happens, the corresponding solution will not produce the maximum or most efficient improvement.

A good example is a busy Wi-Fi hotspot in a coffee shop. Sluggish Wi-Fi is a common complaint when there are more than a few customers in the shop. Often the adopted solution is upgrading either the backhaul speed, the Wi-Fi speed, or both. However, more often than not, the real bottleneck is the server TCP stacks — they throttle data transmissions in reaction to the dynamics in the Wi-Fi network. After upgrading either the backhaul or Wi-Fi speed, the server TCP stacks will increase their throughput simply due to improved available bandwidths in the paths — as a result, the upgrade solutions do provide improvements in the data speed. However, such a solution is not only costly but also hides the real root cause.

Today, numerous sluggish network problems have been misdiagnosed as lacking physical bandwidth or faster physical devices. While costly bandwidth upgrades or device upgrades do appear to solve the problem, the real root cause was never discovered.

An important, yet often ignored and unidentified root cause, is the TCP bottleneck.

TCP is the most important transport layer protocol. Cisco VNI [1] predicts that by 2021, 78% of all Internet traffic will be video. The vast majority of videos distributed over the Internet are transported over TCP, while 100% of the Web traffic is also transported over TCP. Therefore, TCP being the bottleneck is a problem that both consumers and businesses cannot ignore — it is a problem to be reckoned with and solved.

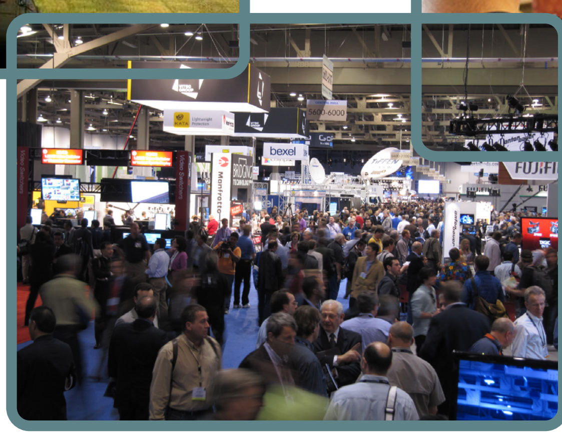

TCP Bottleneck - Hidden and Pervasive

In an ideal world, TCP should never be the bottleneck in any applications — the reason is that TCP has been designed to detect the available bandwidth in its path, and it is designed to transmit data close to the measured available bandwidth.

However, as TCP is also designed to work under all network conditions and configurations, TCP uses a simplified feedback control, allowing it to be universally deployable in today's heterogeneous networks. The flexibility comes with a price — the simplified feedback control prevents TCP from measuring the available bandwidth promptly and accurately.

Consequently, as the measured bandwidth is never accurate or timely, TCP's transmission control becomes conservative — false alarms (falsely detecting a lower bandwidth) are tolerated while over-estimating the real bandwidth is not. The result is that TCP tends to under-utilize the real available bandwidth.

The bandwidth under-utilization problem is magnified when the path includes a wireless link. The issue is that the TCP sender has no way of determining the real cause of a packet loss. If the loss is due to congestion in a middle box, the TCP sender should reduce its transmission rate. However, if the loss is due to random interference in RF transmission, the TCP sender should hold its transmission rate. Therefore, TCP tends to under utilize the available bandwidth whenever the path goes through a noisy wireless link — such is the case when the link is a Wi-Fi or cellular link in a busy environment.

40 years ago, when TCP was invented, most paths did not include a wireless link. Today, a wireless link is ubiquitously included in the last hop to the end user.

TCP bandwidth under-utilization is a pervasive and serious problem, often hidden from users, enterprises, and providers.

How Much can TCP Throughput Drop?

With simplified feedback and conservative rate control, TCP throughput is sensitive to the operating condition. It is well known that TCP throughput can drop exponentially[2] in reaction to increasing RTT (round-trip time) and packet loss rate:

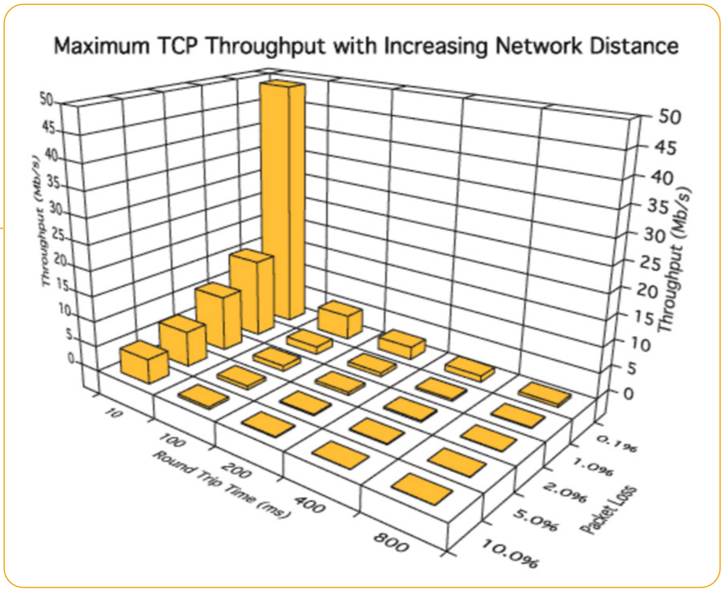

This plot provides the maximum TCP throughput based on a steady-state average analysis. Based on this plot, experts configure their networks and provision physical resources to achieve their design objectives. Even though the average TCP throughput can drop significantly, most networks are configured to ensure that TCP throughput is not the bottleneck in normal operating conditions.

In real environments, both RTT and packet losses vary dynamically. Even when the network is properly designed and configured, dynamic fluctuations in the operating condition for TCP is a fact of life.

Today, RTT varies dynamically due to buffering in the middle boxes and processing delay in the end hosts, while packet loss rate varies dynamically whenever there is a wireless link subjected to radio interference or fading.

Therefore, as the network condition experiences dynamic fluctuations, TCP throughput can drop drastically. From the above plot, even in the average sense, the worse throughput drop can reach 70%-99%. When that happens, the network falls into a congestion collapse.

Don't be misled into thinking that TCP throughput collapses are rare — they often occur in busy wireless networks. In fact, they routinely occur in busy Wi-Fi hotspots (airport, hotel, coffee shops, etc.) and overloaded LTE sites. In a Wi-Fi network, a congestion collapse can take place when more than 5-10 active users are added. This is a problem experienced by almost all users around the world.

Google agrees that TCP is an issue in today's network[3]: "Today TCP's loss-based congestion control—even with the current best of breed, CUBIC11—is the primary cause of these problems."

Misdiagnosis of TCP Bottlenecks

The busy Wi-Fi performance problem is well known — IEEE[4] recognizes this problem; yet they suggest Layer-1 or Layer-2 solutions. IEEE, the world's largest association of technical professionals, has yet to identify TCP as a root cause for poor Wi-Fi performance.

Today, enterprise-grade Wi-Fi AP vendors offer numerous bandwidth solutions using a combination of Layer-1, Layer-2, Layer-3, and even Layer-7 techniques. However, no vendors have offered a solution to directly solve the TCP bottleneck problem.

A similar phenomenon is also observed in the WAN Optimization market. According to Gartner[5], the vendors offer all types of solutions based on inline optimization, caching, compression, deduplication, traffic shaping, etc. However, no vendors[6] have yet offered a solution to directly solve the TCP bottleneck problem.

A reason for not treating the TCP bottleneck problem is that the impact of network dynamics on TCP has not been fully appreciated. While most experts attribute TCP's lower throughput to long RTT, few have paid attention to RTT variance, which is dynamic variation in RTT. This is an issue that has been ignored by all TCP optimization vendors in the industry — Badu is the only vendor that has focused on TCP problems caused by RTT variance.

Video Explosion

Only a few researchers have studied[7] TCP's slowdown problem in reaction to end-to-end delay variance:

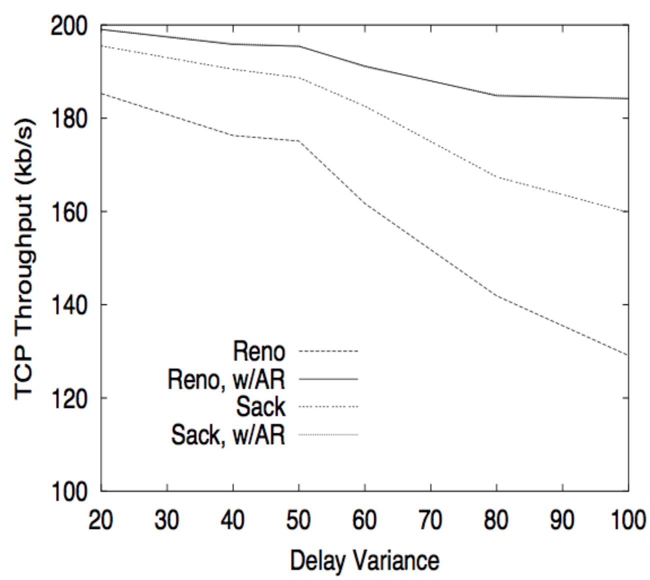

Review

The drop in TCP throughput in reaction to RTT variance has been independently verified in both the field and lab. The following plot is the result of a Badu lab test utilizing a Wi-Fi AP.

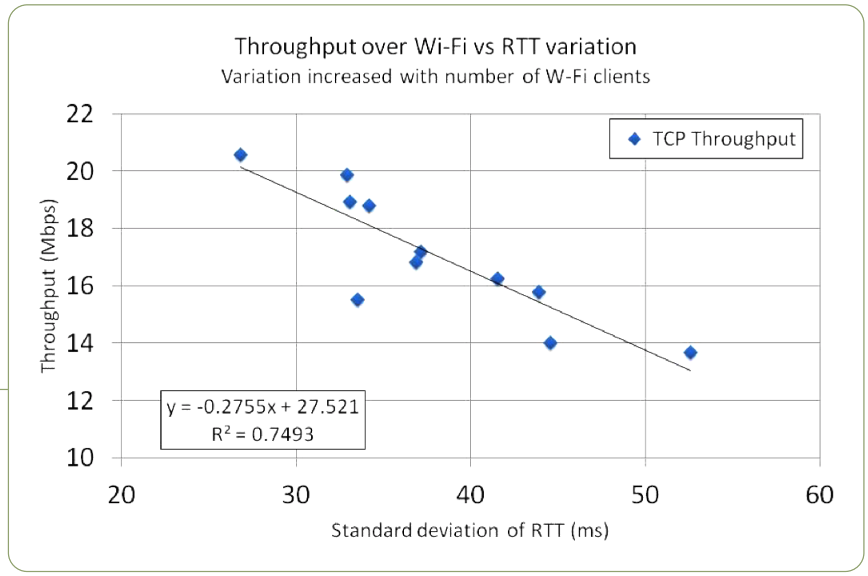

Benefits of TCP De-Bottleneck

Researchers have identified TCP's under-utilization performance problem in wireless links for more than 20 years. Yet, no vendors have brought to market an effective tool to combat TCP wireless woes, until Badu.

Badu's *Warp*TCP solutions have been designed from the ground up to de-bottleneck TCP.

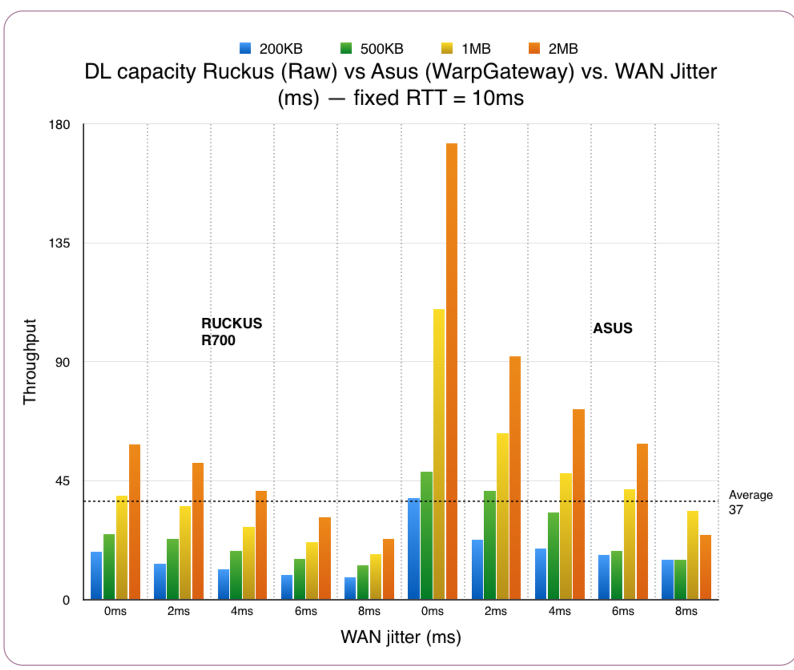

This is the throughput improvement by putting a *Warp*TCP proxy in front of a consumer AP, as compared to an enterprise AP when the WAN connection is jittery. This shows the benefits of TCP de-bottleneck without paying for a backhaul upgrade or an AP upgrade. In the plot, Asus is the consumer AP, while Ruckus is the enterprise AP.

This is the throughput improvement by putting a *Warp*TCP proxy in front of a server with a handset connecting to an LTE base station. This shows the benefits of TCP de-bottleneck without paying for an LTE upgrade.

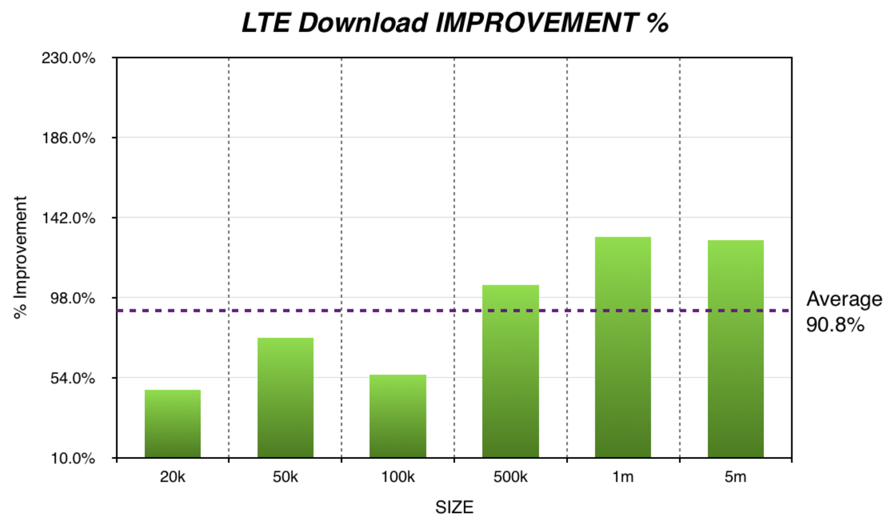

Conclusion

It is time to take action to discover if TCP bottleneck is the root cause of your network's sluggish performance. TCP de-bottleneck can yield the largest improvement and efficiency, without paying for link or device upgrade.

To request a demo, visit: **www.BADUnetworks.com**.

References

1. http://www.cisco.com/c/en/us/solutions/collateral/service--provider/visual-networking-index-vni/mobile-white-paper-c11-520862.htm

2. http://downloads.asperasoft.com/en/technology/shortcomings_of_TCP_2/the_shortcomings_of_TCP_file_transfer_2

3. https://research.google.com/pubs/pub45646.html

4. http://spectrum.ieee.org/telecom/wireless/why-wifi-stinksand-how-to-fix-it

5. http://www.dssi.pt/documentos/fabricantes/riverbed/geral/2016-Magic-Quadrant-Report-on-WAN-Technologies.pdf

6. While vendors such as Riverbed and SilverPeak provide TCP optimization solutions, TCP optimization is only a small part of their solution approach, showing that they don't treat TCP as the root-cause bottleneck.

7. Chan & Ramjee, "TCP/IP Performance over 3G Wireless Links with Rate and Delay Variation" MOBICOM'02

BADU networks

949-310-5390
888-958-7697
info@badunetworks.com

2640 Main Street Irvine CA 92614
http://www.BADUnetworks.com