



**Before the
FEDERAL COMMUNICATIONS COMMISSION
Washington, D.C. 20554**

In the Matter of)	
)	
Request for comments on petition for declaratory ruling regarding Internet management policies)	WC Docket No. 07-52
)	

**COMMENTS OF
THE INFORMATION TECHNOLOGY AND INNOVATION FOUNDATION**

I. Introduction

The Information Technology and Innovation Foundation (ITIF) welcomes the opportunity to comment on the critical issues raised in the Free Press and Vuze complaints against Comcast network management practices. We have formally commented on this matter before but we would like to add a more in-depth explanation on network latency and jitter and the effect of P2P traffic on a broadband connection. This discussion is especially relevant given the fact that Comcast is now partnering with Vonage to ensure good quality for VoIP applications.

The ITIF is a non-profit, non-partisan think tank committed to articulating and advancing a pro-productivity, pro-innovation and pro-technology public policy agenda. Consistent with this mission, ITIF believes that access to high-speed broadband is critical for all citizens if they are to fully participate in and benefit from the increasingly digital economy.

In its “Petition for Declaratory Ruling,” Free Press, et al. asks the Commission to clarify that the “reasonable network management” exception in the Commission’s “Policy Statement on the Appropriate Framework for Broadband Access to the Internet over Wireline Facilities” does not apply to an Internet Service Provider (ISP) – specifically Comcast – “intentionally degrading an application or class of applications.” A key part of this issue relates to the issue of how to best manage latency and jitter on packet switched networks.

The key reasoning behind Comcast's adoption of active network management technology is to improve the performance of real-time applications like VoIP (Voice over IP) and online gaming. Before the implementation of Comcast's current controversial Sandvine protocol-specific network management system, network latency and jitter had become very problematic during peak congestion times. Bandwidth hungry and persistent applications like P2P (Peer to Peer) were not only consuming disproportionate amounts of bandwidth; they were causing congestion storms on the network. This paper explains the basics of latency and jitter and why certain applications like P2P cause excessive amounts of network latency and jitter.

II. Defining latency and Jitter

Latency is a simple measurement of delay. On a computer network, it is the time it takes a bit or packet to traverse a network before arriving at its intended destination and it's generally measured in milliseconds (ms) where 1,000 ms equals one second. The typical latency from the east coast to the west coast on the Internet is approximately 40 milliseconds. However, common latency metrics such as "ping" are quoted as round-trip time which is double the one-way latency so the round-trip ping from the east coast to the west coast is approximately 80 ms.

Jitter is the measure in the variation of delay and it's sort of a more granular version of latency. High jitter conditions are essentially micro-congestion storms that occur at the millisecond level because of packet bursting. If a network fluctuates between 80 and 85 ms of delay, then the jitter is low at 5 ms. But if a network mostly has delays of 20 ms but occasionally spikes to 220 ms, then the jitter is high at 200 ms. While the latter example may have better average delay, its volatility and high jitter characteristics make it less desirable for real-time applications than a network with higher average delay but lower jitter. Real-time applications are very sensitive to these micro-congestion storms while file transfer applications are relatively unaffected.

It should be noted that the word "jitter" is a very technical term that is used infrequently. The word "latency" is commonly used to describe jitter or latency but jitter is the technically correct term to describe the typical problems associated with packet switched networks.

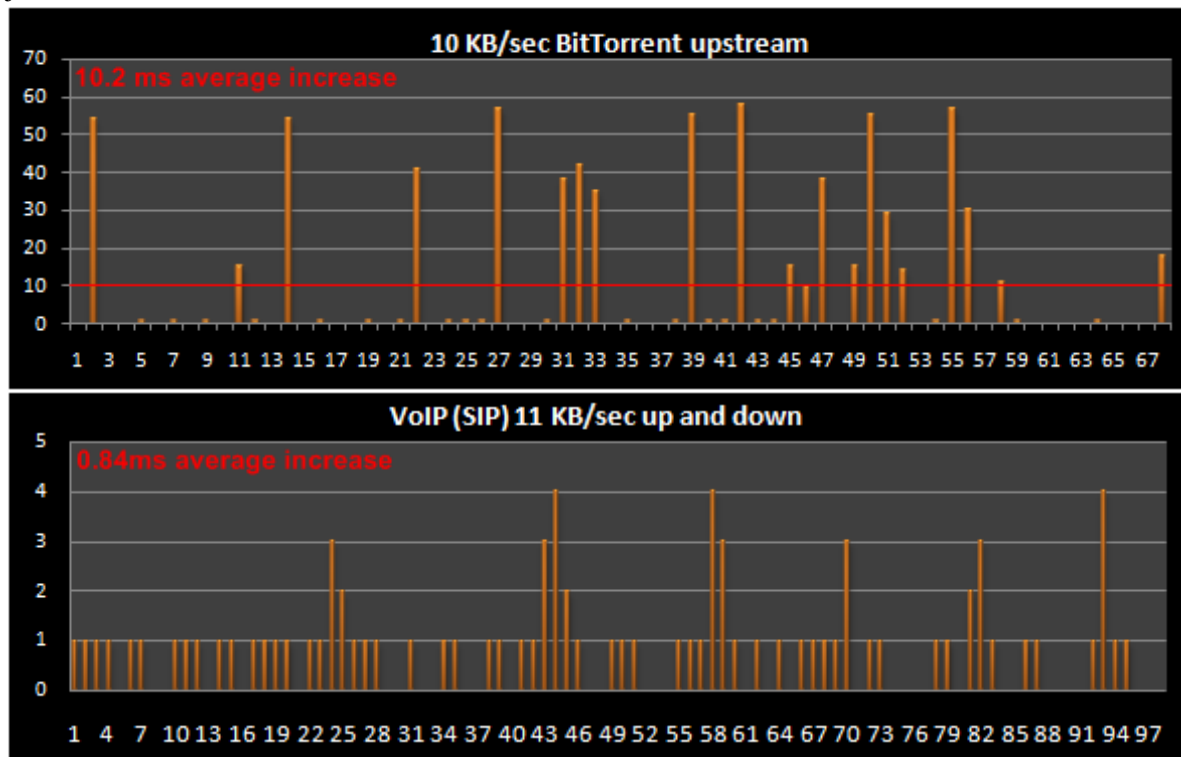
Real-time applications use "isochronous" data transfer which is a way of transmitting data with smooth periodic bursts of data. Online gaming and VoIP for example transmit and receive small bursts of data at 30 and 50 times a second, but the duration of the bursts are very short and these applications use very little total bandwidth. The challenge for real-time isochronous applications is high latency or high jitter because they need their packets delivered with minimum delay and minimum deviation of the periodic delivery schedule.

Under good network conditions, packet switched networks can perform just as well for real-time applications as circuit switched networks with minimal latency and minimal jitter. But under certain load conditions, packet switched networks can have very high jitter. Higher network load conditions generally create higher jitter but it is quite possible for low network load conditions to create jitter as well though less frequently.

III. Network jitter and delay can't be solved by capacity

There are times when relatively low-bandwidth applications can induce high network delay while high-bandwidth applications have little to no effect on an IP network. For example, even two high-bandwidth 1000 KB/sec (kilobyte per second) IPTV real-time video streams over an fiber-to-the-node (FTTN) broadband network create little to no additional jitter but a low-speed P2P file transfer can cause delay to spike up.

The following two charts illustrate the amount of delay induced by BitTorrent (a P2P application) and a VoIP call with similar bandwidth requirements. Despite the fact that the VoIP call is using more bandwidth in both directions, it produces negligible delay. But minimal usage of BitTorrent operating at less than 20% upstream load of a broadband connection can create substantial spikes in delay and high jitter.




IV. Why applications like P2P create high jitter

What's actually causing this is the fact that P2P applications tend to burst out large number of packets at once. This is another side effect of the multi-flow aspect of P2P. The red packets on the right represent P2P packets clogging the network for 85 ms at a time which is a relatively long period of time for VoIP to wait because it needs to send data out every 20 ms.

Under extreme circumstances, it is possible for P2P applications to cause over 1000 ms latency on the downstream side of a broadband connection.

Even if a P2P application is capped at a small fraction of the total capacity by the user, it can still occasionally cause jitter on the network. Even if it's only one out of every 5 VoIP packets being delayed and 4 VoIP packets aren't affected at all, a 20% packet loss is unacceptable so the VoIP application is forced to buffer 85 ms which causes the entire VoIP session to be delayed and that becomes noticeable to people trying to have a real-time conversation.

Real-time isochronous applications on the other hand space out their packets allowing other applications to slip their packets in to the gaps. Real-time isochronous applications don't add any noticeable jitter to the networks they use.



V. Basic packet switched network queuing theory

There are some who claim that quality of service (QoS) doesn't work on the Internet because it's impractical to control every single hop on the Internet. And because the system is unlikely to be perfect, it therefore "doesn't work" and this form of network intelligence should not exist. This is one of the key justifications for net neutrality legislation that bans QoS technology. But this argument is misguided because it argues that only perfect solutions should exist, even though the vast majority of the jitter problems are actually on the last mile of the Internet and they can be dealt with very effectively.

The home network where distances are measured in meters rather than kilometers will always be orders of magnitude faster than the broadband connection because it's much cheaper to build faster networks that are short-haul. Even when we eventually have gigabit upstream connections, our home networks will operate at 10, 40, or 100 gigabits. This large speed mismatch is ripe for large amounts of upstream jitter on the broadband connection. The core interconnects of the Internet must always be orders of magnitudes faster than the edges of the network to aggregate traffic from thousands of user. This large mismatch in speed is by design because it minimizes network congestion but it is ripe for large amounts of downstream jitter on the broadband connection. The up and down speed mismatch and bottleneck on the broadband network is the perfect storm for large amounts of network delay and jitter.

Packets traveling from a faster network to a slower network always have the potential to hit a queuing delay. This is just like cars merging from a 5-lane freeway on to a 3-lane freeway always have the potential of hitting a traffic pile-up. Even when there is sufficient bandwidth to support all the packets, there can be delays when clumps of packets show up at the same time.


Networks with dissimilar speeds are connected by routers or switches with packet queues that temporarily hold packets that can't be transmitted as quickly on the slower network as they arrive on the faster network. Because computer applications tend to burst out large number of packets and even more so with P2P applications, it's possible to have queuing delay even under relatively light P2P loads on networks wherever there are large mismatches in speeds.

VI. Solving jitter problems with QoS

QoS Network Prioritization (a.k.a packet scheduling) schemes can fix jitter problems by rearranging the packets and moving the real-time packets to the front of the transmit queue in the network equipment. Packet scheduling has little to no effect on P2P applications or other non real-time applications because their overall rate of throughput doesn't change.

Without QoS, the user is forced to suspend their P2P application if they want to use real-time applications like VoIP or online gaming or tolerate a decline in quality on their phone call or game. In situations where more than one family member or roommates are sharing a broadband connection, P2P usage is a huge source of friction between housemates. Under a dumb unmanaged network, one housemate would be forced to stop using P2P or the other housemate would be forced to endure a bad performance but it doesn't need to be this way with an intelligent network. Under a QoS optimized network, the P2P user can consume the lion's share of bandwidth while the VoIP or online gamer can get low-bandwidth with minimal delay and this is the best of all worlds.

Implementing a complete QoS solution requires participation on both ends of a broadband link. Ideally, the problem on a DSL broadband network is dealt with on the DSL modem in the home and the DSLAM (Digital Subscriber Line Access Multiplexer) because that's where the queues build up. On a cable broadband network, the problem is ideally dealt with on the cable modem in the home and the CMTS (Cable Modem Termination System) because that's where the queues build up.



A partial solution can be implemented on the consumer end in the home router but it can't completely eliminate the downstream jitter problem because that's ideally handled in the DSLAM or CMTS where the packets queue up. The other problem is that the vast majority of consumers do not purchase expensive high-end home routers with effective QoS technology.

VII. User controlled prioritization

Any broadband user who has ever had the misfortune of living with a P2P user understands the sheer misery P2P applications can inflict. In this situation, one roommate either stops using P2P or the other roommate stops playing their low-bandwidth latency-sensitive online game. Since the P2P application only wants raw bandwidth and the gaming application only wants low latency/jitter and not so much bandwidth, both types of applications should be able to perform optimally if the network can prioritize the gaming packets. While the upstream congestion problem can be handled with a good QoS-enabled home router, the downstream latency and jitter problem is best handled by the ISP's equipment.

One solution to this problem is that the ISP might offer a QoS service where the consumer had access to a web-based control panel that allows them to pick and choose what applications or companies they want to prioritize on their own downstream. A granular web-control interface for end-users may turn out to be too complex for the average consumer and/or too complex to implement for the ISP so a simple opt-in or opt-out on real-time application prioritization might be the better way to go. The consumer already gets to "discriminate" against Apple in favor of Microsoft if they choose an XBOX 360 over an Apple TV so public policy should not bar ISPs from prioritizing applications or services especially when the consumer wants this type of prioritization.

Maybe it's best for the ISP to bundle this service in to their standard service tier or perhaps it's better to offer this service as a premium service for a nominal fee. Some people would love to have the free downstream prioritization service but others may not want to use it or subsidize the service for those who do. There will always be some level of service bundling and it's usually best to let the market place determine where these lines should be drawn rather than have the government declare fees on QoS technology illegal.

VIII. Comcast's protocol agnostic solution and partnership with Vonage

Comcast will be switching to a “protocol agnostic” network management solution by year end but it is important to understand that “protocol agnostic” does not mean that all protocols will be treated equally. Different applications and protocols have fundamentally different requirements for optimum performance and the network needs to accommodate every application and protocol simultaneously as much as possible. What it means is that the network management system will no longer make assumptions about a protocol's bandwidth characteristics even though it's mostly accurate. The new system will attempt to fairly distribute bandwidth amongst users instead of amongst protocols so that it can be completely accurate and fair.

Fairness however does not mean that all protocols should be throttled equally by the network management system, especially when it's intra-user throttling or throttling between applications being used by the same broadband account. Imagine a user who simultaneously uses P2P and VoIP and they've manually reserved about 100 Kbps to their VoIP application by capping their P2P speeds to 2900 Kbps out of a 3000 Kbps connection. But during rare congested times, the ISP can only fairly allocate 2000 Kbps of bandwidth to that users. Does that mean the P2P application should get throttled to 1933 Kbps while the VoIP application gets throttled to 67 Kbps? If we wrongly assumed that protocol agnostic means all applications get equally throttled, we would be making a grave error since throttling a VoIP stream to 67 Kbps when it needs a minimum of 87 Kbps would simply break the VoIP application completely and the phone call is effectively blocked.

The VoIP application fundamentally “throttles” itself to 87 Kbps and allows the P2P application to have 30 times more bandwidth. If that advantage given to P2P is cut to 20 times more bandwidth because the network management scheme reserved the minimum bandwidth for VoIP traffic, would that be a bad kind of discrimination against P2P applications in favor of VoIP? When we understand the characteristics of real-time applications versus the characteristics of P2P applications, the question becomes silly because reserving the bandwidth for VoIP or any other similar real-time application is the obvious correct thing to do because we don't want it to break. This is precisely why Vonage (a VoIP telephone service provider that competes with Comcast) is keenly interested in partnering with Comcast to ensure that the new network management scheme provides safety mechanisms for VoIP traffic.

At this point, it isn't clear if Comcast intends to offer more advanced forms of network management that actively reshuffles the packets to reduce network jitter for real-time applications but this should be encouraged rather than be seen as something evil. It is important to understand that fixing jitter for real-time applications does not adversely affect the throughput of file transfer applications like P2P. Network management isn't just about bandwidth management; it's about good traffic engineering that ensures all applications work well together and that all users get a good experience.

When we think about the fact that Comcast is effectively helping out one of their telephony competitors to provide better Internet service for their customers, we realize that there is nothing sinister about favoring one application over another. In fact, it would be nice if Comcast extended this kind of basic protection for all low-bandwidth real-time applications like online gaming and other VoIP applications like Skype.

IX. Conclusion

ITIF appreciates the opportunity to comment on the issues raised by the Commission in its “Request for Comments”. We are happy to see that Comcast is going the extra mile to ensure that their new protocol-agnostic policy doesn’t adversely affect real-time applications. The Commission will be making a new ruling soon on the matter of Comcast network management and we hope that this paper is helpful in guiding public policy.

Respectfully submitted,
George Ou

Senior Analyst
Information Technology and Innovation Foundation
202-360-4964