Evaluation of "A Fairer, Faster Internet Protocol" through Network Simulator 2 (NS2)

Kyurim Rhee
April 30-2009

***Abstract*— According to Bob Briscoe, the author of "A Fairer, Faster Internet, TCP- the way we share bandwidth needs a makeover"[1], the current method of utilizing the channel through bandwidth allocation is not optimal at all. Different users of the internet request and require different bandwidth. Most users have small files to transfer, but some users exchange huge amount of data. To prevent those with large data transfers from hogging the network resources, some ISP's have implemented throttling to keep the bandwidth open for the "normal" users. In this article, Briscoe argues that with a new method of handling the traffic, there needs to be no need for throttling. In fact, the "normal" users who are simply browsing the internet or checking their email will be able to get their files even faster, and those who are transferring large data are not delayed in any way. All this is done without the ISP mediating via throttling. In this paper, the current network model and the ISP throttled model will be examined. The proposed model will also be examined.**

***Index Terms*— NS-2, NS2, A Fairer, Faster Internet Protocol, Bob Briscoe, Queueing, Packet Scheduling, Throttling, WRR.**

I. Introduction

The current internet protocol is not suited for a fair trafficking of packets. It gives unfair advantage to those who utilize the network on a constant basis, yet gives little share of the resources to those who use it on small bursts. A good example is a comparison of normal user and heavy user. A typical normal user will check their email or browse the web which creates small bursts of traffic. A heavy user on the other hand may be one who downloads large files which take long time to transfer. According to the article "Faster, Fairer Internet Protocol" by Bob Briscoe, a typical neighborhood is composed of user ratio of heavy user to light user by 1:4. However, the traffic created by the light users are typically only active during 5% of the time where as the traffic of the heavy users may be active all the time as they may be using some type of file sharing program which may run unattended[1]. This creates an unfair situation as the 20% users (heavy users) are using a significant amount of the network resources and are causing the rest of the user's (light users) data transfer to become slow. To address this problem, some internet service providers have started to "throttle" the traffic of the heavy users. This solution may not be the best solution as will be described in section II of this paper. The paper will go into some details of possibly addressing this issue that has been proposed by Bob Briscoe in his article.

This paper is divided into four sections. In section I, we introduce the IEEE published paper and its main ideas. Section II. will discuss the motivation for a new internet protocol and evaluate the current network setups and its short comings by evaluating the typical neighborhood topology with both a normal traffic scenario and that of an ISP throttled case. These cases will be evaluated by a simulation using NS2 network simulator [2] to gather quantitative results to compare the different cases. In section III, we will review how we can improve upon the current network by discussing the method introduced in the IEEE published article "A Fairer, Faster Internet Protocol". In section IV, we conclude the paper by reiterating the key points.

II. Simulation: To Throttle or Not to Throttle?

A. User Usage in Time

To model the typical neighborhood scenario, we first discuss the user traffic behaviors. As the ratio of heavy user to normal user is 1:4, 1 heavy user and 4 normal user traffic was created. The heavy users will always create traffic, hence we have modeled them as the always being active as shown in figure 1. The normal users will create an exponentially randomly distributed inter-arrival time with the users being active only 5% of the time. However, to be able to distinctively observe when the active times of the normal users are in order to observe the corresponding results, I have opted to manually set the random start/stop times of the normal user as shown on figure 1 rows 2~4. Figure 1 shows the times of a user's data transfer times for the NS2 simulation for the duration of 60 seconds.

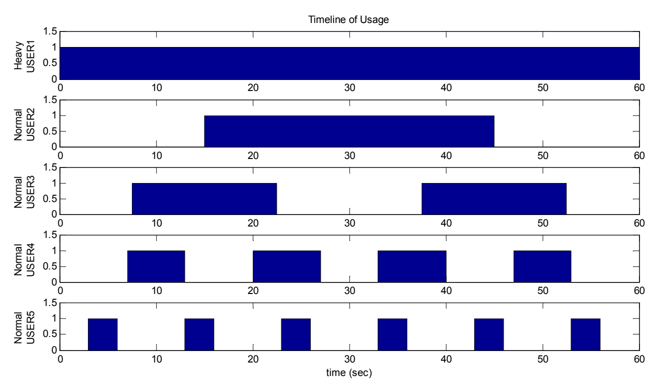

Fig1. Time line of User Usage

B. Topology of Typical Network Model and the ISP Throttled Model

The typical neighborhood network can be modeled in a dumbbell topology of each end nodes representing an end user of different types shown in figure 2. For creating the simulation, one node was created to represent the heavy user and the four other nodes were created to represent the reset of the normal users. The right side of the dumbbell are the corresponding servers that each end users are sending data to.

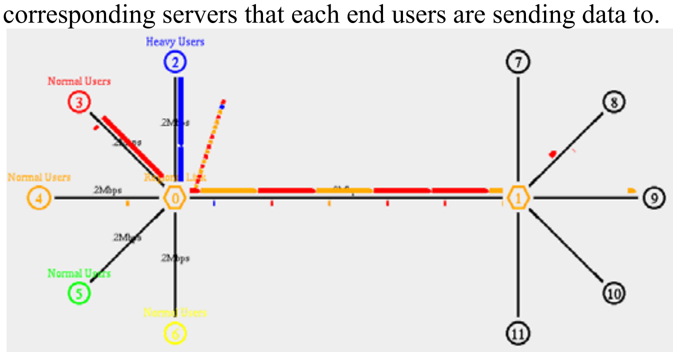

Fig2. Topology of a network (Non-throttled)

There are various ways ISP's can throttle each users. For the simulation, we have picked a simple case of limiting the data rate of the link between the heavy user and the regional link from 2Mbps to 10% of its original capacity of 0.2Mbps. This topology is shown in figure 3. Note the red "throttled" link between Heavy user and the Regional Link.

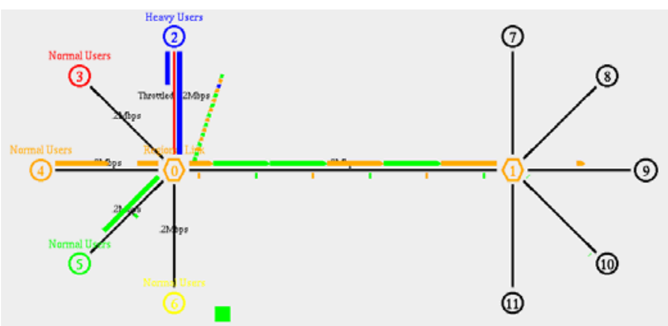

Fig3. Topology of a network (Throttled)

C. TCP Window Size

TCP window size was observed in order to detect when and how many packets are being sent at a time. This data can also be use to determine when the packets are being dropped due to an overflow in the regional link. Figures 4 and 5 shows the window size of non-throttled and throttled traffic, respectively. The bottom graph of each respective graph represents the sum of the window sizes of all nodes in the neighborhood network. Note that when the sum of all window size peaks, there are too many packets at the regional link, hence the packets getting dropped for all active users at that time, i.e., at t = 13sec, t = 37sec. We can also observe that the throttling will decrease the rate at which the window size is increased. When throttled, a packet is sent at a slower rate at the sender's link, which slows down the cycle of increasing the window size.

D. Throughput

The throughput is a direct measure of who is utilizing the channel. By observing this with respect to time, we can observe who is utilizing the regional link and at what times. Figures 6 and 7 shows the throughput of non-throttled and throttled traffic, respectively. The bottom graph of each respective graph represents the sum of the throughput of all users. This in turn represents the throughput of the regional link. Note that for throttled user of figure 7, the data interval becomes larger showing the effect of throttling.

Using the acquired data, the average throughput was calculated as shown in figure 8 and 9 for cases of non-throttled and throttled.

E. Simulation Results

Tables 1 and 2 summarizes the simulation results.

	(MBps)	(%)
Non-Throttled	Avg. Throughput	Channel Utilization
1 Heavy User	0.10302	34
4 Normal User	0.14823	49
Regional Link	0.25125	84
Throttled	Avg. Throughput	Channel Usage
1 Heavy User	0.01692	6
4 Normal User	0.19301	64
Regional Link	0.20993	70

Table 1. Utilization of channel

Gain in Performance (throttle/non-throttle)	
1 Heavy User	0.16
4 Normal User	1.30
Regional Link	0.84

Table 2. Performance gain by throttling

Results of the NS2 simulations show that by throttling, the four normal users will gain an additional 30% increase in the network performance. However, the one heavy user will see a drastic drop in their performance by 84%. Moreover, the regional link will also suffer an overall performance by 16%.

Indeed the ISP throttling can be used to limit the overall usage of the heavy users and give back some bandwidth to the normal users. However, the drastic decrease in performance experienced by the heavy users may not justify the 30% increase in performance by the normal users. From a network engineer's point of view, it is also less attractive to see a 16% decrease in overall network performance.

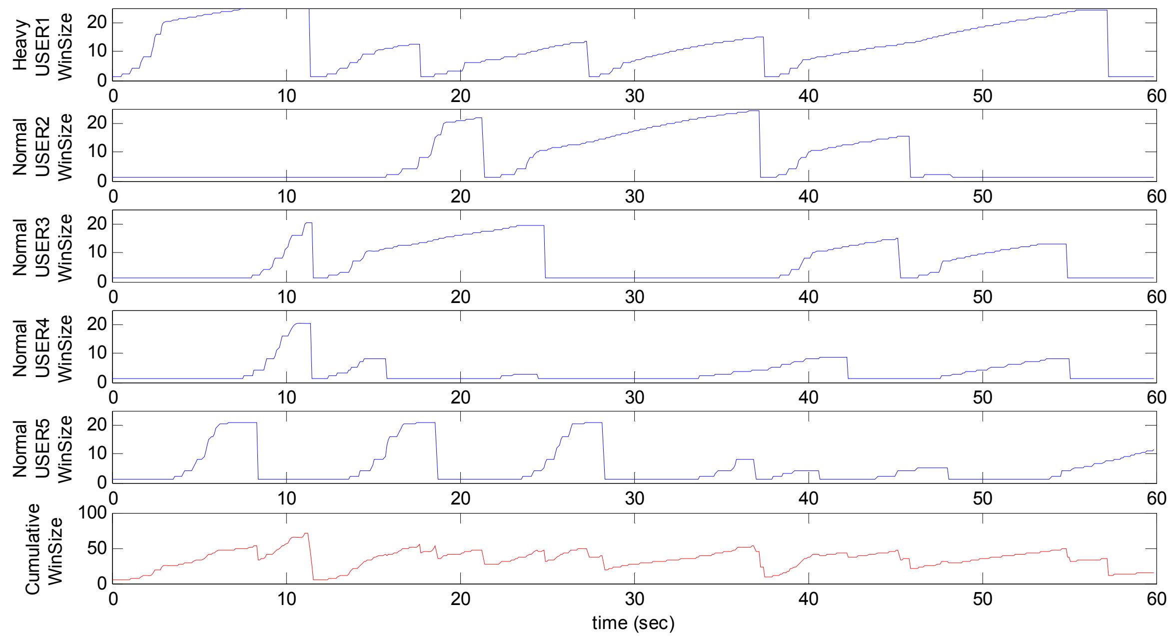

Fig4. TCP window size of user (Non-throttled)

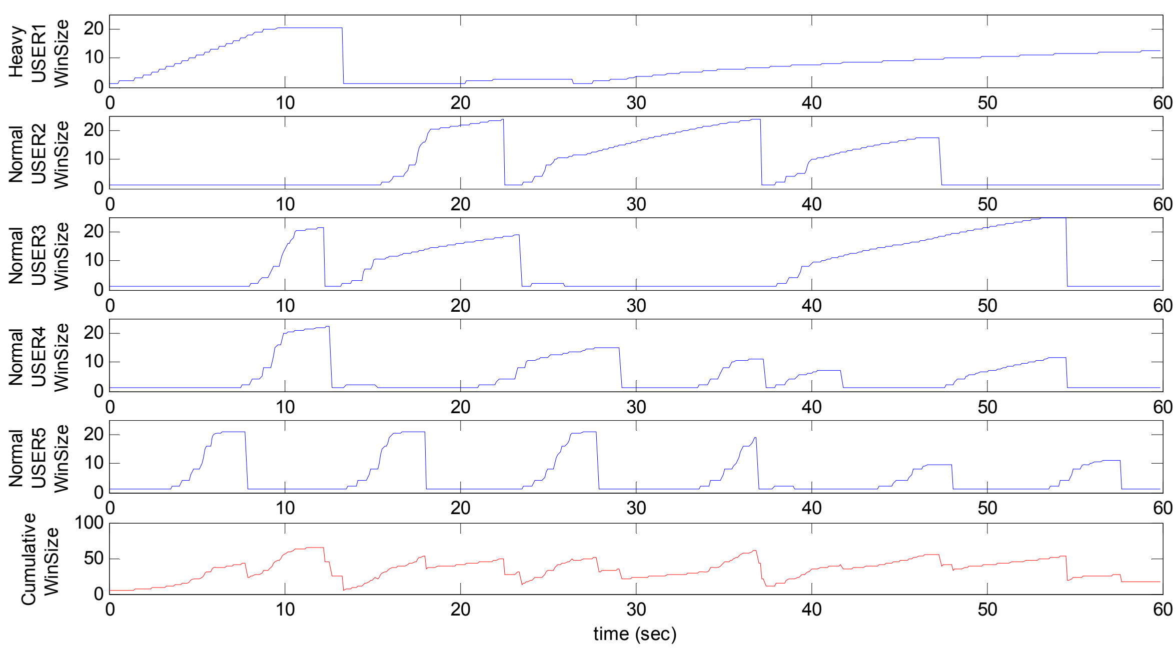

Fig5. TCP window size of user (Throttled USER1)

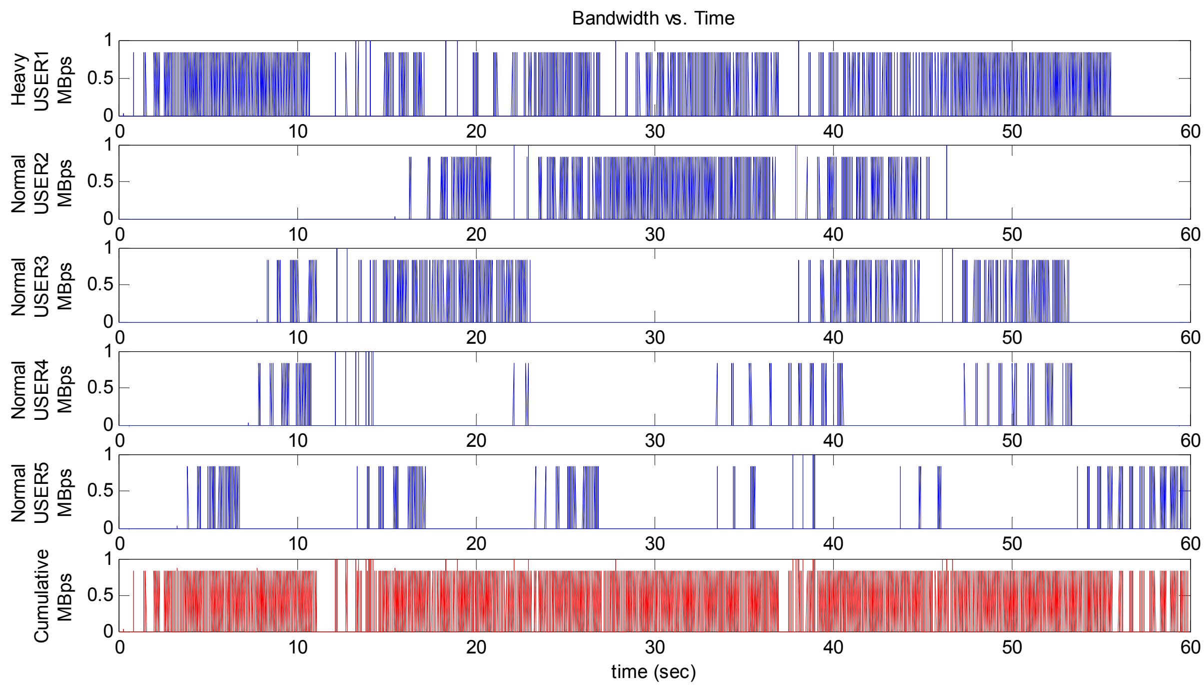

Fig6. Throughput of user (non-throttled)

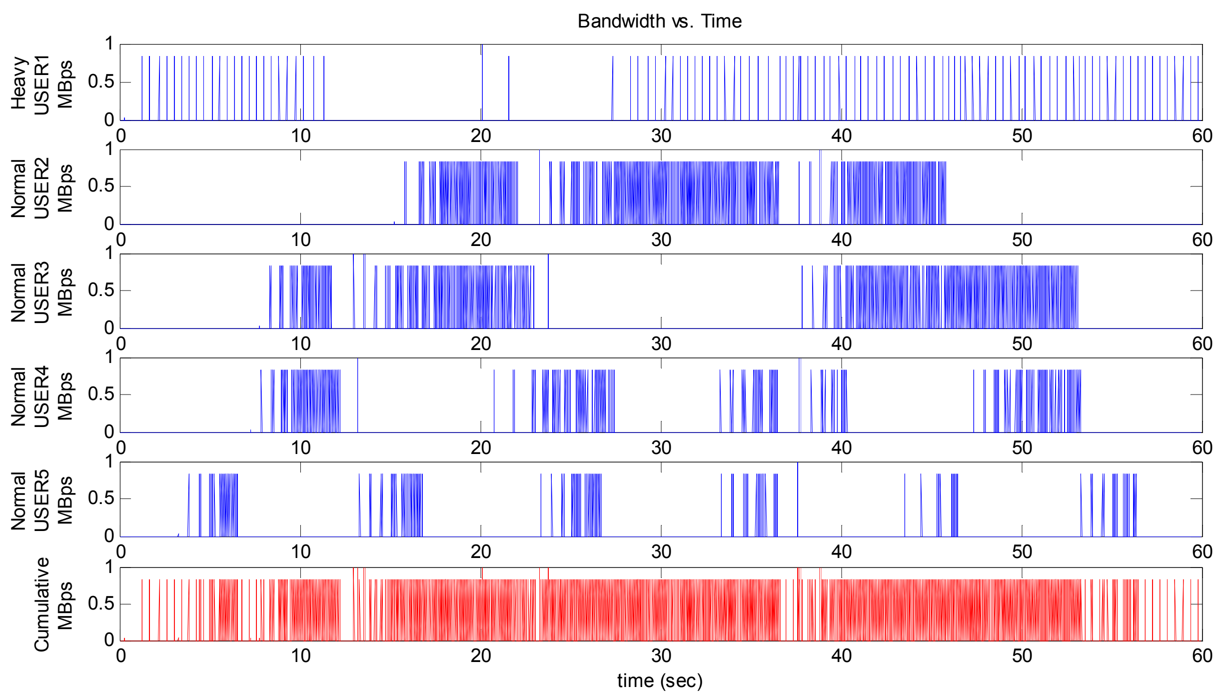

Fig7. Throughput of user (Throttled USER1)

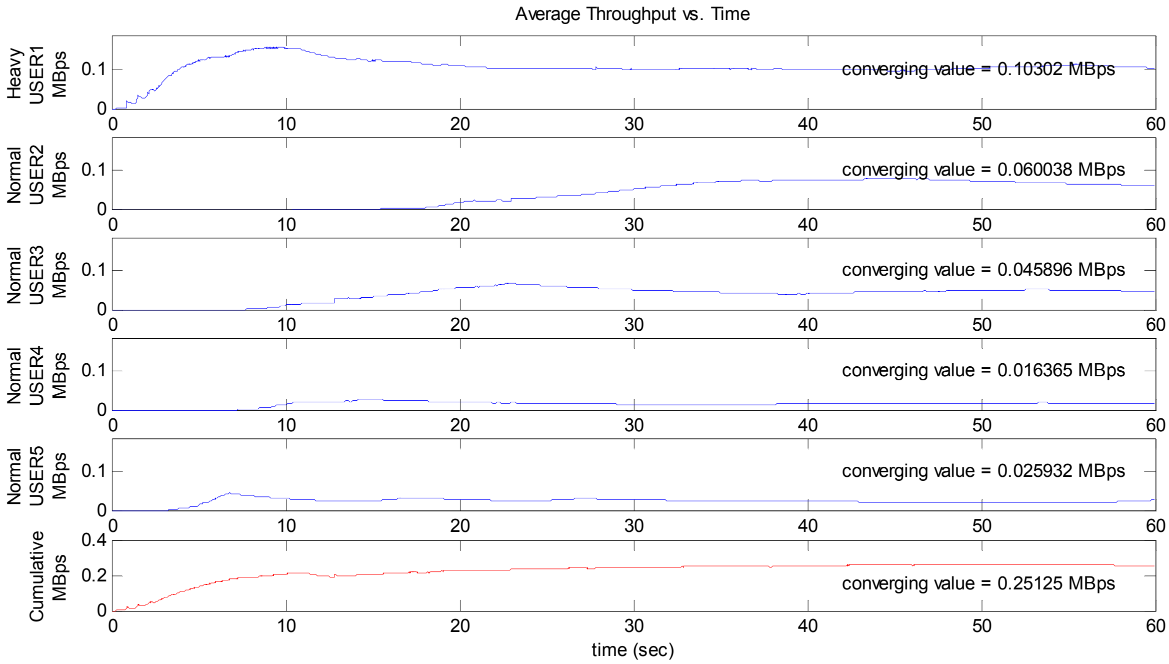

Fig8. Average throughput of a user (non-throttled)

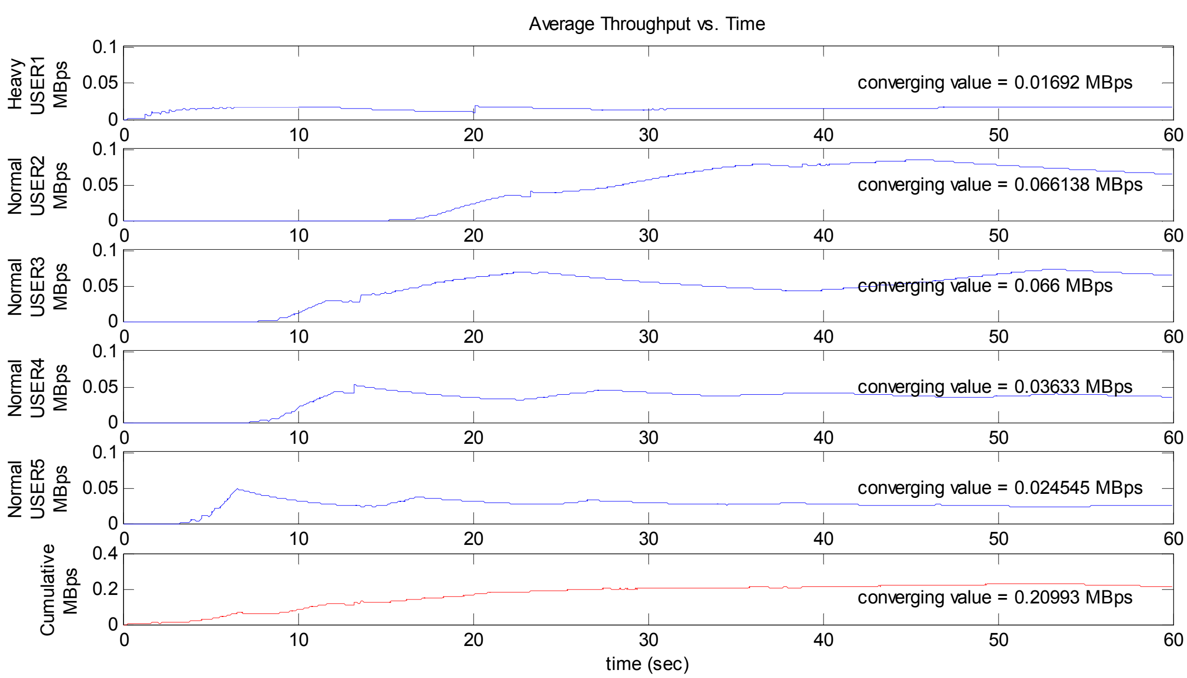

Fig9. Average throughput of a user (Throttled USER1)

III. Better Solution: Proposed Internet Protocol

A. *Proposed Network Model*

The Weighed TCP sharing method mentioned in "A Fairer, Faster Internet Protocol"[1] is an alternative solution to ISP throttling. A weighting system can be created by setting weight to each packet with respect to the size of the total file. This is accomplished by setting a weighing system with small file transfer requests to high weight and large file transfer requests to low weight. When a packet with high weight is requested, then the channel will allocate a large portion of the bandwidth to this packet. Once all the high weighing packets have been transferred, the rest of the bandwidth can be allocated to the low weighing transfers. This model can be implemented using priority queueing or WRR (weighted round robin) scheduling method[3] shown in figure 10. When a high priority queue is occupied, the rate at which low priority queue empties will decrease. This will create an overflow in the low priority queue resulting in low priority packets to reduce their TCP window size and automatically "backing off" their transmission.

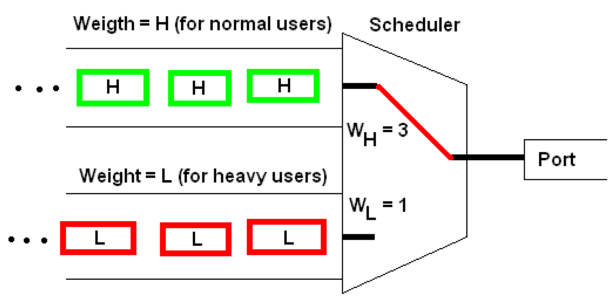

Fig10. WRR scheduler

B. *Incentive to lower the weight by each user*

The question we now ask is, why doesn't everybody set their weight to high? How do we make the lower priority packets set their packet weight to low? According to Bob Briscoe, a user will have a monthly allocation of the total number of high priority weight. If the user sets all his data transfers to high priority and does not use their high priority weights sparingly, then at the end of each month he would have run out of high priority weights and all his data transfers will be set to lower weight. This will result in their data transfer rate to be low regardless of the priority of their packets. The idea of this system is to be sparing with the priority high weight in order to create a fair trafficking between high and low priority packets.

C. *Results*

By implementing a weighted round robin scheduling or priority queueing system, it is true that one can build a fair and faster network where a heavy user will slow down their transfers rate and allow the normal user to get their data fast and efficiently. This weighted scheduling system will also utilize the capacity of a network to the fullest.

IV. Conclusion

Through the experiment results, we can see that ISP throttling may not be the best solution to the problem of heavy users over utilizing the resources. Bob Briscoe with his article on "Faster, Fairer Internet Protocol" addresses the issue and proposes a solution. However, some may argue that creating a monthly "point system" for the high priority weight packets seems somewhat comparable to ISP's putting a cap on the amount of data to be transferred for the month. Upon closer analysis, as this system doesn't completely stop the data transfer when the point limit is met, this solution, though not as elegant as one may have hoped, seems plausible solution to this issue.

Acknowledgment

The author would like to thank Dr. Brian Mark for his discussions on the issue.

References

[1] B. Briscoe, "A Fairer, Faster Internet Protocol", IEEE Spectrum, Dec 2008 http://www.spectrum.ieee.org/dec08/7027/1

[2] NS2 – Network Simulator 2 http://www.isi.edu/nsnam/ns/

[3] B. Mark, "Packet Processing and Scheduling" http://teal.gmu.edu/~bmark/742Sp09/slides/PDF/sched.pdf

[4] SFQ (stochastic fairness queueing) http://www.opalsoft.net/qos/DS-25.htm