

# Investigation into TCP Congestion Control Performance for a Wireless ISP

J. Lillis<sup>1</sup>, L. Guan<sup>1</sup>, X. G. Wang<sup>2</sup>, A. Grigg<sup>3</sup>, W. Dargie<sup>4</sup>

<sup>1</sup>Department of Computer Science, Loughborough University, LE11 3TU, UK

<sup>2</sup>School of Computing, Communication and Electronics, University of Plymouth, Plymouth, UK

<sup>3</sup>Systems Engineering Innovation Centre (SEIC), Loughborough University, LE11 3TU, UK

<sup>4</sup>Department of Computer Networks, Technical University of Dresden, Germany

{Corresponding author: L.Guan@lboro.ac.uk}

## Abstract

Wireless technology and devices are becoming ever more pervasive and embedded in our lives. In particular, wireless access to the internet is defining the way that we learn, work and socialize. The primary motivation for undertaking this investigation was the author's involvement in a community broadband project. The contributions of the paper have real impacts for industry and the results of this investigation have been used to provide recommendations to a community wireless internet provider and highlight the wider impact on the UK broadband network. Five major contributions have been achieved: (1) Results from both wired and wireless scenarios are critically evaluated and then concluded; (2) Accurate simulations of 802.11b network in NS2 are validated with live network tests; (3) Evidence is given through both simulation and real network tests to show that current TCP standards are inefficient on wireless networks; (4) A new TCP congestion control algorithm is proposed as well as an outline for a fresh approach; (5) Industry reactions are then given to recommendations for network changes as a result of this investigation.

**Key Words:** TCP Congestion Control, Performance and QoS, Wired and Wireless Networks

## 1. Introduction

### 1.1 Motivation

Wireless access to the internet has become increasingly pervasive and is now embedded in many devices – mobile phones, PDAs, media centres, laptops and MP3 players are just some common examples. Despite the overwhelming trend of internet users migrating to some form of wireless connection, standard TCP implementations have not been reviewed.

The author's long-term involvement in a rural community broadband project, FramBroadband, provided the impetus to investigate how performance on

the wireless network could be improved. FramBroadband was founded in 2003, aiming to supply broadband via wireless technologies, due to extremely slow deployment of ADSL in the small town of Framlingham and the surrounding rural villages located in Suffolk, UK. The network is formed of a number of wireless base stations, interconnected via wireless point-to-point links. The gateway to the internet is provided from the Technology Centre in Framlingham; an office building with fibre internet backhaul (c.f. Fig. 1). Although ADSL is now available in Framlingham and the surrounding villages, many users experience severe performance issues in rural areas and, in numerous cases, businesses and residents are located too far from the telephone exchange to receive a reliable service at adequate speed, or indeed any level of service at all. The initial aim was to supply an alternative to ADSL but it is now hoped to provide services superior to ADSL. The wireless infrastructure at time of writing has capacity to deliver up to 20Mbps, pending an upgrade of fibre connections to the internet. At the present time FramBroadband supplies up to 12Mbps.

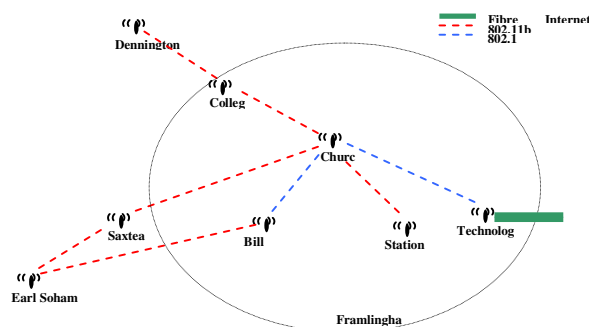


Figure 1 : FramBroadband Network Topology of Base Stations

Although wireless internet service providers are sparse in the UK at the moment, there is evidence to suggest this will change in the future. The major issue facing non-fibre connected areas is 'the last mile' of an internet connection – the telephone wire from the exchange to user premises. Factors such as line loss and distance from the exchange limit availability, performance and

reliability. For services above 2Mbps, a distance within around 3.5km is required [1]. This is telephone line distance; therefore premises in urban areas that connect via an indirect cable route to the exchange are affected. ADSL Max/2+ [2] attempted to solve this problem by increasing the range and giving users their maximum possible line speed up to 8Mbps. However, on average, an 'up to 8Mbps' broadband connection in the UK was found to deliver only 2.7Mbps [3]. Although primary infrastructure owner BT so far has failed to provide a solution to the limitations of ADSL, they have recognised the challenge and have been conducting research into WiMax as an alternative [4]. This provides encouraging evidence that wireless technologies will play a crucial role in coming years, and supports the wider value of this project.

## 1.2 Supporting Research with FramBroadband

Congestion control [5] is a serious issue for FramBroadband, especially over bottlenecks created by point-to-point links connecting network segments. Congestion can also be created by wireless link performance, which is the result of intermittent factors such as rain, wind and even humidity.

A large operational challenge for FramBroadband is web browsing performance over the long distance point-to-point links. Therefore, the focus was to identify potential for TCP performance improvements for existing users and provide recommendations for how FramBroadband may address TCP congestion control on its network. The project would further aim to link the context of results to the future role of wireless providing the 'last mile' of broadband delivery. This project outcomes would also be directly applicable to all users of wired broadband who connect through a home wireless access point.

A comparison of TCP on wired and wireless networks was required. Many papers investigated alternative TCP versions on wired networks, but there was little research into a comprehensive comparison of their performance in a wired and wireless environment. TCP was originally developed as a congestion control protocol for wired networks and makes the assumption that all packet loss is due to congestion. Therefore this project aimed to show that this assumption was not transferable to a wireless environment, and results in an inefficient utilization of available bandwidth.

A commonly used TCP version, NewReno, performs congestion control on the basis of packet loss. An alternative version, Vegas [6], uses round trip times instead [7]. Veno [8], which was specifically developed with wireless networks in mind was subsequently added

to the investigation project. It was also agreed with FramBroadband, that results from a simulation replicating a portion of their network would be most valuable.

The paper continues with a review of network simulators, congestion control and wireless medium access control. The simulation was developed by writing several NS2 scenarios in a wired environment, before moving on to a wireless domain. Results from both wired and wireless scenarios are then critically evaluated. Final recommendations are then given for FramBroadband, along with the implications of these findings and suggestions for future research.

## 2. Simulation Investigation Results

### 2.1 Wired Scenario Network Setup

Figure 2 shows the topology used for wired network scenarios. Two UDP sources with a Pareto traffic distribution were used to create congestion over the bottleneck between nodes 3 and 4 during the FTP transfer. The Pareto distribution provides a more suitable traffic model as it incorporates traffic burstiness. This allowed comparison of TCP Vegas and NewReno with a greater degree of realism.




Figure 2 : Wired Scenario Network Topology

In a wired network environment, with a number of randomly generated traffic sources, Vegas was shown to have enhanced performance when compared to NewReno, in particular, maintaining a higher level of throughput but with a lower and more constant round trip time.

Figure 3 and Table 1 show that Vegas achieves higher throughput than NewReno, with the differentiating factor being the avoidance of packet loss. The Pareto traffic sources showed that Vegas was also able to react to changing congestion levels, rather than just converging to a static window size when faced with a constant rate of background traffic.

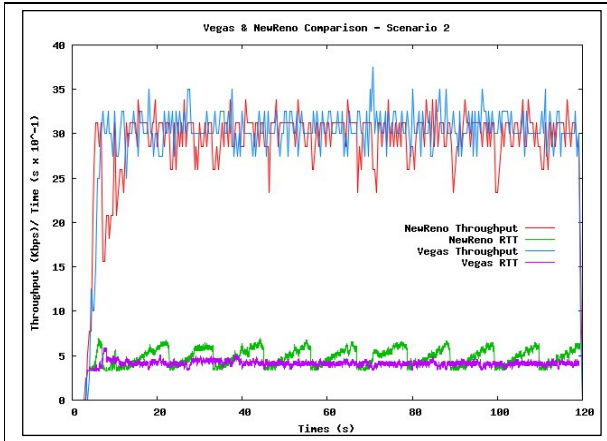


Figure 3 : Performance in Wired Scenario

	Reno	Vegas
<b>TX Packets</b>	3317	3491
<b>RX Packets</b>	3294	3491
<b>Dropped Packets</b>	23	0
<b>Drop Ratio</b>	0.693	0
<b>Throughput (Bytes)</b>	3294000	3491000
<b>Average Delay (ms)</b>	337	262
<b>Average RTT (ms)</b>	488	413

Table 1 : Performance in Wired Scenario

The wired scenario was then extended to include a loss model to evaluate how the two protocol versions performed in an environment with a high level of loss. Due to Vegas' fundamentally different approach to congestion control, based on round trip time instead of loss, the performance of the protocol when presented with random loss was investigated. It was shown that Vegas had potential for substantial performance increase over NewReno, with a nearly 70% higher throughput.

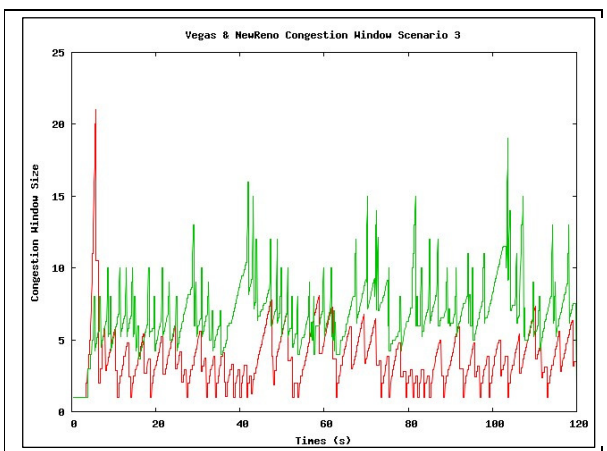


Figure 4 Congestion Window Comparison

	NewReno	Vegas
TX Packets	1066	1740
RX Packets	922	1668
Dropped Packets	74	72
Drop Ratio	6.942	4.138
Throughput (Bytes)	992000	1668000
Average Delay (ms)	196	194
Average RTT (ms)	347	345

Table 2 : Congestion Window Performance

The loss model introduced a very high loss rate for a wired network of 5% on the last link of the TCP route. Table 2 shows that in terms of throughput Vegas achieved much higher performance; however the delay for the connection is approximately the same. NewReno is forced to reduce its sending rate due to periodic loss, rather than being allowed to fill up buffers. Figure 4 shows that the difference in throughput was clearly due to Vegas maintaining a consistently larger congestion window. This behaviour can be explained by the difference in Vegas' loss recovery scheme – because Vegas keeps a much more accurate record of RTT it is able to respond to the losses caused by the error model much quicker than NewReno. It is also possible that NewReno could experience multiple losses at a given congestion window size, especially due to its approach of expanding window size until loss occurs, combined with the error model introduced. In the event of multiple loss, again NewReno is disadvantaged by the potential of the congestion window being reduced multiple times for loss at a particular rate. The error model scenario gives strong evidence that Vegas could outperform NewReno in environments of high loss. This provided justification to continue the comparison in a wireless environment.

## 2.2 Wireless Simulation

Before results generation took place, research was conducted into NS2's ability to model signal loss. Signal loss has a high impact on throughput and random loss over a wireless network. A simulation was written to evaluate two propagation models in NS2 over a varying distance between source and receiver. The 'Shadowing' model was shown to produce seemingly realistic degradation. The commonly used TwoRayGround, did not use a realistic random probability function, and cut off at a threshold value rather than a gradual decline.

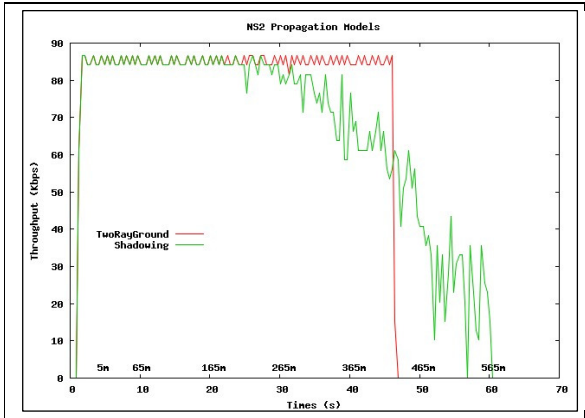


Figure 5 : Signal Loss Test

A portion of the FramBroadband network in Saxtead was then modelled in NS2. Attention was paid to ensure that the simulation replicated the real life setup of the network. This included examining how to specify the physical properties of equipment used, such as antenna gain. Positions of a number of current users were approximated using mapping, then replicated in the topology area.

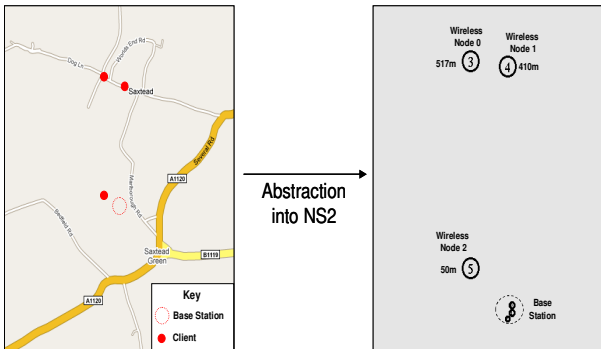


Figure 6 : Topology of Wireless Simulation

The 802.11b standard was configured in NS2 and then validated against a live network test and found to produce throughput to within 10kbytes/s.

TCP Veno was then added to NS2, with the aim of confirming increased performance claims by the developers on wireless networks. However, testing of this protocol showed negligible performance increase, highlighting inaccurate modelling in NS2 of random packet loss in a wireless environment. Although NS2 appeared to model signal fading realistically, this was at the physical level and was based on the power generated by the reception of each wireless frame. This manifested as a drop in available bandwidth at the transport layer, without actually modelling random errors/loss (supported by wider research). Writing a random error model for NS2 was not in scope of the project, but it was

key to show that standard implementations of TCP could be inefficient on wireless networks. This was achieved by migrating the effects of noisy wireless conditions on packet loss to a wired simulation, as well as the characteristics of the FramBroadband network such as bandwidth.

The loss model was then modified in the following ways:

- (1) Link bandwidth was set to 5meg, approximately that of an optimum B link; delays were updated to be more realistic inline with FramBroadband network;
- (2) Fu [8] claims throughput increase of up to 80%, on a typical wireless network with 1% random bit error and hence the error model was updated from 5% to 1%;
- (3) TCP Sinks were changed to SACK1 to support Veno implementation.

	NewReno	Vegas	Veno
<b>TX Packets</b>	23432	34192	37696
<b>RX Packets</b>	23156	33833	37301
<b>Lost Packets</b>	276	359	395
<b>Loss Ratio</b>	1.178	1.050	1.048
<b>Throughput (Bytes)</b>	23156000	33833000	37301000
<b>Average Delay (ms)</b>	31	30	30
<b>Average RTT (ms)</b>	56	54	56

Table 3 : Performance Comparison in Wireless Test

In Table 3, Veno was shown to outperform NewReno in a variety of performance metrics, notably throughput. The results from the modified scenario were much more in line with the performance expectations of Veno. With a 1% random loss, Veno performs over 60% better than NewReno in terms of throughput – with a marginally lower average delay. Veno also outperforms Vegas, in terms of throughput. Perhaps the most important issue this scenario highlighted was the fact that NewReno, the most common implementation of TCP was beaten by both Vegas and Veno in all performance metrics. This scenario also appeared to validate the shortcoming in NS2 with regards to modelling random wireless bit errors. The error model essentially replicated random loss experienced in a wireless environment, so although this scenario is on a wired network, it is justifiable to assume these trends would be transferable.

### 3. Live Testing

To optimize the results of this investigation, it was desirable to test alternative TCP versions on the network. Research was conducted into compiling

swappable algorithms into the Linux operating system kernel, and the methodology provided to FramBroadband. This was also used to perform preliminary tests on the network that showed potential benefits of alternative TCP versions, as illustrated by Table 4.

Test Number	NewReno (kbps)	Vegas (kbps)
1	436	460
2	322	280
3	484	560
4	419	487
5	504	593

Table 4 : Live Network Testing

#### 4. Proposed VenO Modification

This paper has shown the need to review current implementations of TCP algorithms for wireless networks, as well as a possible review of the approach of congestion control. The following recommendations for a new TCP algorithm are shown below.

Proposed VenO Modifications
<ul style="list-style-type: none"> <li>• Change slow start algorithm so that it expands congestion window exponentially until actual throughput drops below expected throughput.</li> <li>• At the point where the actual throughput drops below expected, initiate VenO's congestion avoidance algorithm, but at a much slower rate of expansion, for example expand the window every four acknowledged packets.</li> <li>• Modify recovery mechanism – if loss is encountered during uncongested state, reduce window by 1/5. If loss is encountered during congestion, initially reduce the window by 1/4, then by another 1/4 if congestion persists.</li> </ul>

#### 5. Industry Reaction

FramBroadband have reacted to recommendations for network changes as a result of this investigation, and are investigating migrating backhaul links to different frequency bands to commonly used 2.4GHz. This was suggested to reduce the random loss that would adversely affect standard TCP implementations, resulting in dropped web pages and poor FTP throughput. FramBroadband have also recognized their potential role as a test bed for new wireless network technologies and importantly as a proof of concept for the future direction of the UK broadband network.

#### 6. Conclusions and Future Work

This paper has presented an argument for a completely different approach to TCP congestion control and provided substantial performance comparison based on simulation and live network testing results. The evidence provided has real impacts for industry. It shows the need to develop congestion control for wireless networks; enabling more effective and reliable communication that will lead to both increased and cheaper levels of service. This paper has also highlighted congestion control as a scalable challenge for provision of broadband over wireless. In future work, the outlined modification to TCP with advanced features in congestion control protocol fairness on wireless networks could be implemented.

#### 7. References

[1] Thinkbroadband, 2008, Technical FAQ, available at <http://www.thinkbroadband.com/faq/sections/technical.html#166>

[2] Riezenman, M.J, 2003, Extending broadband's reach, IEEE Spectrum, Vol. 40, No. 3, p20-21.

[3] Times Online, 2007a, Broadband speeds 'slower than advertised', Times Newspapers Ltd., available at <http://ieeexplore.ieee.org/search/wrapper.jsp?arnumber=1184890>

[4] Times Online, 2007b, BT raises prospect of entry into Wi-Max technology, TimesNewspapers Ltd., available at <http://ieeexplore.ieee.org/search/wrapper.jsp?arnumber=1184890>

[5] Welzl, M., Network Congestion Control: Managing Internet Traffic, John Wiley & Sons. 2005.

[6] Brakmo, L.S, Peterson, L.L, 1995, TCP Vegas: End to End Congestion Avoidance on a Global Internet, IEEE Journal on selected areas in communications, Vol. 13, No 8, 1465 – 1480.

[7] Kurata, K., Hasegawa, G., Murata, M., Fairness Comparisons Between TCP Reno and TCP Vegas for Future Deployment of TCP Vegas, 2007. Osaka University Japan, available at [http://www.isoc.org/inet2000/cdproceedings/2d/2d\\_2.htm](http://www.isoc.org/inet2000/cdproceedings/2d/2d_2.htm)

[8] Fu, C.P., Liew, S.C, 2003, TCP VenO: TCP Enhancement for Transmission Over Wireless Access Networks, IEEE Journal on Selected Areas in Communications, Vol. 24, No.2, p216 – 228.