

Modelling TCP Dynamics in Wireless Networks

D.J. Leith, P.Clifford
Hamilton Institute, NUI Maynooth

Abstract—In this paper we develop an analytic model of the behaviour of competing TCP flows in wireless networks. A key feature of this work is that we take explicit account of the interactions between competing TCP flows. This allows us to study issues such as responsiveness and the impact of the flow AIMD parameters on fairness.

I. INTRODUCTION

Network traffic is currently dominated by data traffic (web, email, media downloads, etc.) carried via the TCP reliable transport protocol and this situation is likely to continue for some time. While there have been many empirical studies of TCP performance in wireless networks, analytic results are much less common. In this paper we develop an analytic model of the behaviour of competing TCP flows in wireless networks. A key feature of this work that distinguishes it from, for example, the well known square-root formula of Padhye et al [12] is that we take explicit account of the interactions between competing TCP flows. This allows us to study issues such as responsiveness and the impact of the flow AIMD parameters on fairness. While our model is applicable in a quite general wireless setting, in this paper we focus on illustrating its predictive power in the context of 802.11e networks.

II. MODELLING TCP DYNAMICS

We consider N wireless stations competing for access to a shared wireless medium and with station $s \in [1, N]$, being the source of n_s TCP flows. The TCP flow destinations may be either wired or wireless stations (we allow flows to have arbitrary round-trip times) but it is assumed that the paths of all flows include a shared wireless hop.

A TCP sender maintains a state variable $cwnd$ that determines the number of sent but unacknowledged data packets; that is, packets in flight. TCP operates an AIMD congestion avoidance strategy, whereby a TCP sender increases its $cwnd$ by α packets each round-trip time until a packet loss is detected, at which point $cwnd$ is backed-off to $\beta \times cwnd$ and the process starts again. We have that packet losses occur either due to overflow (as a result of increases in $cwnd$) of the queue at the bottleneck link or due to a loss on the wireless medium.

We make the following initial assumptions:

- (i) The wireless hop is the bottleneck link for all flows (this assumption is relaxed later). TCP data packets are therefore queued at the wireless station interface queues.
- (ii) The station interface queues are sized such that they do not empty when a TCP flow backs off its congestion window $cwnd$ on detecting congestion. Hence, the wireless stations are saturated (always have a data packet to send) and the activity on the wireless medium is thereby decoupled from the values of the individual flow $cwnd$ s.




Fig. 1. Evolution of window size over a congestion epoch. $T(k)$ is the length of the congestion epoch in seconds.

- (ii) From the viewpoint of station s , the wireless link provides a stochastic service rate with mean value $E[B_s]$ packets/sec. The link suffers from random packet losses, with loss probability $p_{s,loss}$ per packet transmitted. By (ii), we can assume that the service and loss rates are statistically independent of the flow $cwnd$ s.

Since the bottleneck queue for a flow is the wireless station interface queue, we have that TCP flows from the same wireless station share a common interface queue but flows at different stations are held in separate queues. Note that we already have, by (ii), that when modelling TCP $cwnd$ evolution we can consider each wireless station separately.

Consider a station s and TCP flow i . The evolution of the $cwnd$ of a typical flow as a function of time, over the k th congestion epoch, is depicted in Figure 1. We denote by $w_{s,i}(k)$ the $cwnd$ value of flow i just before backoff and let the nominal AIMD increase and decrease parameters be, respectively, $\alpha_{s,i}$ and $\beta_{s,i}$. In addition, we denote by $t_{s,ai}(k)$ the time at which the number of packets in flight belonging to flow i is equal to $\beta_{s,i}w_{s,i}(k)$, by $t_{s,bi}(k)$ the time at which packet loss occurs, and $t_{s,ci}(k)$ is the time at which the flow is informed of this. The evolution of $cwnd$ does not evolve linearly with time due to (i) the stochastic nature of the service rate and (ii) the effect of the bottleneck queue filling and the resulting variation in RTT. The RTT of the i th flow increases according to $RTT_{s,i}(t) = T_{d_{s,i}} + (1 + q_s(t))/B_s(t)$ where $T_{d_{s,i}}$ is the propagation delay of the wired leg of flow i 's path, $0 \leq q_s(t) \leq q_{s,max}$ denotes the number of packets in the queue at station s and $B_s(t)$ is the (stochastic) service rate of the station. Note that we do not assume that every source experiences a drop when congestion occurs.

We have from the AIMD rule that

$$w_{s,i}(k+1) = b_{s,i}(k)w_{s,i}(k) + a_{s,i}(k)T_s(k) \quad (1)$$

where $b_{s,i}(k) \in \{\beta_{s,i}, 1\}$ corresponding to whether the flow

experiences a packet loss or not at the k th congestion event, $a_{s,i}(k)$ is the *effective* AIMD increase rate in packets per second and $T_s(k)$ is the duration of the k th congestion epoch.

A. Noise-free channel

Owing to the stochastic service rate and random losses on the wireless channel, $a_{s,i}(k)$, $b_{s,i}(k)$, $T_s(k)$ and thus $w_{s,i}(k)$ are random variables. We consider initially the situation without random packet losses on the wireless channel where $p_{s,loss} = 0$ and packet losses occur only due to queue overflow. On congestion we have by assumption (ii) that the aggregate flow throughput must equal the queue service rate,

$$\sum_{i=1}^{n_s} \frac{w_{s,i}(k)}{RTT_{s,i,max}(k)} = B_s(k) \quad (2)$$

where $B_s(k)$ denotes the effective service rate of station s at time k and $RTT_{s,i,max}(k) = T_{d_{s,i}} + (1 + q_{s,max})/B_s(k)$ is the RTT of flow i at the end of the k th congestion epoch. Also,

$$\begin{aligned} \sum_{i=1}^{n_s} \frac{w_{s,i}(k+1)}{RTT_{s,i,max}(k+1)} &= \sum_{i=1}^{n_s} \frac{b_{s,i}(k)w_{s,i}(k) + a_{s,i}(k)T_s(k)}{RTT_{s,i,max}(k+1)} \\ &= B_s(k+1) \end{aligned}$$

By assumptions (ii) and (iii), we have that probability distribution of $B_s(k)$, and so in particular the mean value of $B_s(k)$, is independent of the congestion epoch k . Letting $\gamma_{s,i} = 1/E[RTT_{s,i,max}]$ and taking expectations therefore yields

$$\sum_{i=1}^{n_s} \gamma_{s,i} E[w_{s,i}(k)] = E[B_s] \quad (3)$$

and

$$\sum_{i=1}^{n_s} E[b_{s,i}] \gamma_{s,i} E[w_{s,i}(k)] + \gamma_{s,i} E[a_{s,i}] E[T_s] = E[B_s] \quad (4)$$

where we have assumed that the probability of a flow backing off at a congestion event is independent of w_i and that the effective AIMD increase parameter is independent of the congestion epoch and its duration¹. Hence,

$$E[T_s] = \frac{1}{\sum_{i=1}^{n_s} \gamma_{s,i} E[a_{s,i}]} \left(\sum_{i=1}^{n_s} (1 - E[b_{s,i}]) \gamma_{s,i} E[w_{s,i}(k)] \right)$$

Inserting this expression into (1) we finally obtain

$$\begin{aligned} \gamma_{s,i} E[w_{s,i}(k+1)] &= E[b_{s,i}] \gamma_{s,i} E[w_{s,i}(k)] \\ &+ \frac{\gamma_{s,i} a_{s,i}}{\sum_{j=1}^{n_s} \gamma_{s,j} E[a_{s,j}]} \left(\sum_{j=1}^{n_s} (1 - E[b_{s,j}]) \gamma_{s,j} E[w_{s,j}(k)] \right) \end{aligned} \quad (5)$$

The dynamics of the collection of flows at wireless station s are thus described by

$$W_s(k+1) = A_s W_s(k) \quad (6)$$

where $W_s(k) = [\gamma_{s,1} E[w_{s,1}(k)] \dots \gamma_{s,n_s} E[w_{s,n_s}(k)]]^T$,

$$A_s = \begin{bmatrix} E[b_{s,1}] & 0 & \dots & 0 \\ 0 & E[b_{s,2}] & 0 & 0 \\ \vdots & 0 & \ddots & 0 \\ 0 & 0 & \dots & E[b_{s,n_s}] \end{bmatrix} + \Gamma_s \begin{bmatrix} \gamma_{s,1} E[a_{s,1}] \\ \gamma_{s,2} E[a_{s,2}] \\ \vdots \\ \gamma_{s,n_s} E[a_{s,n_s}] \end{bmatrix} [1 - E[b_{s,1}], \dots, 1 - E[b_{s,n_s}]]$$

with $\Gamma_s = 1/\sum_{j=1}^{n_s} \gamma_{s,j} E[a_{s,j}]$, and where the initial conditions satisfy $\sum_{i=1}^{n_s} \gamma_{s,i} E[w_{s,i}(0)] = E[B_s]$. The model involves the quantities $\gamma_{s,i}$, $E[a_{s,i}]$ and $E[b_{s,i}]$ which can be obtained as follows. We have that $\gamma_{s,i} = E[RTT_{s,i,max}]$ and $E[RTT_{s,i,max}]$ equals $T_{d_{s,i}} + (q_{s,max} + 1)/E[B_s]$. Letting $\lambda_{s,i}$ denote the probability of flow i backing off when a congestion event occurs, we have that $E[b_{s,i}] = \lambda_{s,i}(1 - \beta_{s,i})$. The value of $\lambda_{s,i}$ can be measured from network traces or set to a value of interest, e.g. $\lambda_{s,i} = 1$ when a flow experiences a packet loss at every congestion event. We typically approximate $E[a_{s,i}]$

by $\gamma_{s,i} \alpha_{s,i}$.

B. Noisy channel

Consider now the situation where random packet losses can occur on the wireless channel. The queue need not therefore be full on congestion. By assumption (ii) we have, however, that the queue is not empty on congestion. Hence, on congestion the aggregate flow throughput $\sum_{i=1}^{n_s} w_{s,i}(k)/RTT_{s,i}(k)$ must equal the queue service rate. $RTT_{s,i}(k)$ is the RTT experienced by the i 'th flow when congestion occurs in the k 'th congestion epoch. Note that $E[RTT_{s,i}]$ is $T_{d_{s,i}} + E[(q_s + 1)/B_s]$, where $E[q_s]$ is the average queue occupancy at which a packet loss occurs, and $E[RTT_{s,i}]$ is equal to $RTT_{s,i,max}(k)$ when congestion occurs due to queue overflow. For low loss rates we have that $E[q_s]$ is close to the maximum queue size $q_{s,max}$, but as the loss rate increases the mean value of q_s will decrease.

C. Properties

Since $E[a_{s,i}] > 0$ and $0 < E[b_{s,i}] < 1$, the A_s matrix in (6) is positive and has similar form (namely, a positive matrix formed from the sum of a diagonal element and a rank one element) to that studied by Shorten et al [8], [7] in the context of purely wired networks. The matrix has positive real eigenvalues, with a unique largest eigenvalue λ_1 of unity. Thus the network dynamics are globally exponentially stable with equilibrium point equal, to within a normalising factor, to the eigenvector of λ_1 . Specifically, the equilibrium value of $w_{s,i}$ of flow i is proportional to $E[a_{s,i}]/(1 - E[b_{s,i}])$. The magnitude of the second largest eigenvalue of A_s bounds the convergence properties of the network. This eigenvalue depends on the backoff factors of the flows. Specifically, it follows from the eigenvalue interlacing theorem in [11] that the second eigenvalue is upper bounded by largest mean backoff factor $E[b_{s,i}]$. The convergence rate of the network mean congestion windows is thus also bounded by the largest mean backoff factor $b_{s,max} = \max_i E[b_{s,i}]$, with for example

¹The validity of these assumptions is assessed below by comparing model predictions against packet level simulations.




Fig. 2. Throughput of competing TCP uploads and downloads (NS simulation, 10 upload TCP flows, 10 download TCP flows, infrastructure mode 802.11b WLAN, TCP SACK).

the 95% rise time bounded by $\log 0.05 / \log b_{s,max}$.

Comment: Wired networks. The foregoing model can also be applied directly to networks with a wired bottleneck link. The service rate is then constant, thus $E[B]=B$.

Comment: Bottleneck location. When the path includes a wireless hop but the bottleneck link is a wired one, then we can no longer assume that the queues at the wireless stations never empty (i.e. that the wireless stations are saturated). The wireless hop essentially acts as a stochastic, lossy delay. Nevertheless, provided the queue at the wired bottleneck is sized such that it does not empty on congestion and the wireless channel is lightly loaded, the wireless delay and losses may be assumed to be approximately independent of the flow congestion windows and the foregoing model remains valid. Similarly, when the flows at some stations have a wireless bottleneck and flows at other stations have a wired bottleneck.

III. TCP DYNAMICS IN 802.11E WLANs

In this section we consider the application of our model in an 802.11e wireless network. Conventional 802.11 WLANs are known to suffer from unfairness between competing TCP flows. This is discussed in Section III-A, together with an approach for using the flexibility provided by the new 802.11e MAC to resolve this problem. In later sections we then verify the accuracy of our model in the 802.11e context.

A. TCP Unfairness in 802.11 Networks

Figure 2 illustrates the behaviour of competing TCP upload and download flows over a conventional 802.11b WLAN. Gross unfairness between the throughput achieved by competing flows is evident. Such behaviour has also been noted previously in empirical studies by [2], [3], [4].

The source of this highly undesirable behaviour is rooted in a negative interaction between the MAC layer contention mechanism and the transport layer congestion control action. The MAC layer enforces per-station fair access to the wireless

channel; that is, n stations competing for access to the wireless channel are each able to secure approximately a $1/n$ share of the total available transmission opportunities [2]. Hence, if we have n wireless stations and one AP, each station (including the AP) is able to gain only a $1/(n+1)$ share of transmission opportunities. By allocating an equal share of packet transmissions to each wireless node, the 802.11 MAC allows $n/(n+1)$ of transmissions to be upload TCP data/download TCP ACKs yet only $1/(n+1)$ (the AP's share of medium access) to be download TCP data/upload TCP ACKs. For larger numbers of stations, n , this MAC layer action leads to substantial forward/reverse path asymmetry at the transport layer.

Asymmetry in the forward and reverse path packet transmission rate is a known source of poor TCP performance in wired networks, e.g. see [1]. Firstly, for TCP uploads path asymmetry can create frequent TCP ACK losses at the AP, in which case a situation can easily occur where a newly started TCP flow loses the ACK packets associated with its first few data transmissions, inducing persistent timeouts. This effect is evident in Figure 2 where it can be seen that a number of upload flows are completely starved of throughput. Secondly, unfairness is created between download flows and upload flows since download flow throughput is constrained by the ability of the AP to win transmission opportunities. This is illustrated for example in Figure 2 where it can be seen that upload flows achieve nearly two orders of magnitude greater throughput than competing download flows.

Following [9], [10], fairness can be restored by using the flexibility provided by the 802.11e MAC as follows.

- (i) Strictly prioritise TCP ACKs at the AP and wireless stations. This can be achieved using the 802.11e *AIFS* and CW_{min} parameters. Allowing TCP ACKs unrestricted access to the wireless channel does not lead to the channel being flooded. Instead, it ensures that the volume of TCP ACKs is regulated by the transport layer rather than the MAC layer. In this way the volume of TCP ACKs will be matched to the volume of TCP data packets, thereby restoring forward/reverse path symmetry at the transport layer.
- (ii) The *TXOP* packet bursting mechanism in 802.11e provides a straightforward and fine grained mechanism for prioritising TCP download data packets at the AP. Since the download TCP data traffic gains a $1/(n_u+1)$ share of transmission opportunities, by transmitting n_d packets (one packet to each of the n_d download destination stations) at each transmission opportunity it can be immediately seen that we restore the $n_d/(n_u+n_d)$ fair share to the TCP download traffic.

Revisiting the example in Figure 2, the impact of the proposed prioritisation approach can be seen in Figure 3. It can be seen that fairness is restored between the competing TCP flows. The 802.11e MAC parameter settings used in this example (with an 11Mbps PHY) for both TCP uploads and downloads are summarised in Table I.

With this approach to selecting the 802.11e MAC parameters, the adverse interaction between the MAC layer contention mechanism and transport layer congestion control is avoided. Assuming that the interface queues are sized large




Fig. 3. Throughput of competing TCP uploads and downloads; first 10 flows are TCP uploads, remainder are TCP downloads. (NS simulation: 802.11e parameters as in Table I).

		$AIFS$ (slots)	CW_{min}	TX_{op} (packets)
AP	Upload ACKS	0	2	1
	Download data	4	32	n_d
wireless station	Download ACKS	0	32	1
	Upload data	4	32	1

TABLE I
802.11E MAC PARAMETERS WITH 11MBS PHY

enough that they do not empty when TCP flows backoff, the wireless stations are saturated (always have a packet to send). Recalling that TCP ACK packets are strictly prioritised in the 802.11e scheme, a very simple approximate model of the MAC behaviour can be obtained if we assume that a TCP data transmission and its corresponding TCP ACK reply occur back to back. In this case the 802.11 model of Bianchi can be applied directly. The accuracy of this approximation is illustrated in Figure 4. We observe excellent agreement between the model and simulation except for very small values of CW_{min} where the collision probability is high (greater than about 0.3, corresponding to more than 30% of packet transmissions failing due to collisions). When the collision probability is high, multiple TCP backoff and timeout events become frequent, violating the assumptions on which our model is based. However, it can be seen from the figures that such high collision probabilities are associated with CW_{min} values less than the 802.11 standard value of 32, and so are of little relevance in the present context.

It follows from the Bianchi model that the service rate received by a station and packet losses are both stochastic but independent of the flow congestion windows and that each station receives the same mean service rate. Hence, the TCP dynamics model in Section II is applicable.

B. RTT Unfairness

We use the topology shown in Figure 5, where the WLAN operates in infrastructure mode. With this topology we can consider a mix of upload and download flows. Further, by varying the bandwidth B of the wired link, the bottleneck in

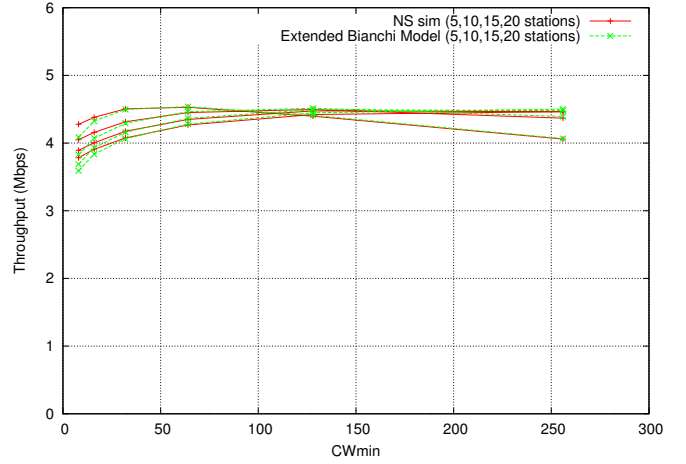


Fig. 4. 802.11e theory versus simulation, varying numbers of upload flows.

the network can be varied between the wired and wireless hops. An 802.11e WLAN with 11Mbps PHY is used, in which case the wired link acts as the bottleneck when its bandwidth B is less than about 5Mbps, whereas the wireless hop acts as the bottleneck for higher values of B .




Fig. 5. Network topology used to compare characteristics of TCP flows with wired/wireless bottleneck link. ($T = 10ms$, $T_1 = 0ms$, $T_2 = 40ms$, wired link queue $0.0001B$, 11Mbps PHY, TCP uploads from stn A to D1 and from stn B to D2, TCP downloads vice versa.)

To start with we compare the behaviour of two upload TCP flows with that of two download TCP flows. When the wireless link is the bottleneck, a key difference is that in the case of TCP uploads the TCP data packets are queued separately at each wireless station, whereas download flows compete via a shared queue at the AP. Access to the wireless channel is regulated by the ability of the wireless stations to secure transmission opportunities for their data packets. The MAC enforces station access independent of the AIMD parameters of the competing TCP flows and hence it can be expected that the bandwidth share achieved by competing TCP upload flows at different stations is invariant with respect to $\alpha_{s,i}$ and $\beta_{s,i}$. This behaviour is confirmed by simulation, see Figure 6. In contrast, download flows compete via a shared queue at the AP and, as noted in Section II-C, in equilibrium the peak congestion windows are proportional to $E[a_{s,i}]/(1 - E[b_{s,i}])$. This behaviour is also illustrated in Figure 6, together with the corresponding model predictions.

An immediate consequence is that wireless upload flows do not suffer from the RTT unfairness that is ubiquitous in wired TCP networks. To demonstrate this behaviour, Figure 7(a) shows simulation results as the bandwidth, B , of the wired bottleneck is varied. Also shown are the corresponding predic-




Fig. 6. Impact on fairness of varying the AIMD parameter increase parameter α of second flow. (NS simulations, two TCP flows, topology as in Figure 5, bandwidth $B=10$ Mbs yields wireless bottleneck).


tions of our analytic model. When the wired link bandwidth is low, the wired link acts as the bottleneck and unfairness exists between the competing TCP upload flows as a result of their different round-trip times². When the wired link bandwidth is increased, thereby shifting the bottleneck to the wireless link, this unfairness disappears. It can be seen that the transition between these regimes is quite abrupt, as might be expected. The predictive power of our analytic model is clear from this figure. Further confirmation of the insensitivity of fairness to RTT when the bottleneck link is the wireless hop is provided in Figure 7(b).

C. Convergence Rate

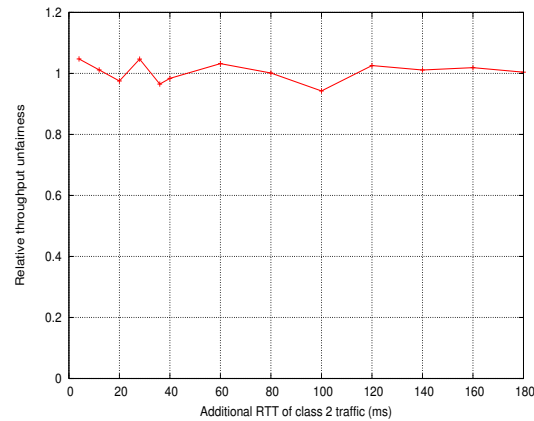
The convergence rate, or responsiveness, of a network of TCP flows is a measure of the time that the network takes to reach steady state following start-up of a new flow or other such disturbance. Once again, we compare the behaviour of TCP uploads and TCP downloads over a wireless link.

When the wireless link is the bottleneck, download TCP data packets for all flows share a common bottleneck queue at the AP, with packet drops largely arising from the aggregate action of the competing TCP flows. The analysis in Section II indicates that the convergence rate measured in congestion epochs is then determined by the AIMD mean backoff factors $E[b_{s,i}]$ of the competing TCP flows, with the convergence time increasing exponentially as the AIMD backoff factor β is increased. Figure 8 shows *cwnd* time histories of TCP download flows where the wireless AP is the bottleneck and thus the flows compete via a common queue. When the backoff factors are all 0.5 (the situation with standard TCP) our model predicts that the 95% convergence time is 4 congestion epochs and it can be seen that this is in good agreement with the results in Figure 8(a). Similarly, when the backoff factors are 0.8, the model predicts a 95% convergence time of 14 congestion epochs, which can again be seen to be in good agreement with Figure 8(b).

²It can be seen that the ratio of throughputs varies with the wired link bandwidth. This is associated with changes in the pattern of packet drops as network conditions are varied - it can be seen that the model predictions, which take account of the packet drop patterns, are in excellent agreement with the measured values.



(a) varying wired link bandwidth B (wired link is network bottleneck when B is less than around 5Mbs and wireless link is bottleneck for larger values).



(b) varying RTT with wireless bottleneck (bandwidth $B=10$ Mbs)

Fig. 7. Impact on fairness of bottleneck link location and RTT. (NS simulation, two TCP uploads, topology as in Figure 5).

In contrast to this behaviour, Figure 9 shows the *cwnd* histories of TCP upload flows from different wireless stations. For TCP uploads the TCP data packets are queued separately at each wireless station. Hence, on startup a new TCP flow will typically not experience any data packet drops until its probing action has led to the interface queue at its own station filling. Convergence following startup of a new flow is therefore largely *independent* of the aggregate action of the network of TCP flows and in this respect is fundamentally different from the download case. Figure 9 illustrates the convergence in a wireless network following the startup of a second TCP upload flow. It can be seen that the new flow increases its congestion window monotonically and experiences no packet drops until its steady state value is reached. The latter is determined by the interface queue size and the delay-bandwidth product of the path. An immediate consequence of this behaviour is that the convergence time measured in congestion epochs in the wireless case is largely insensitive to the AIMD backoff parameter β , see Figure 9. The convergence time measured in seconds is, of course, still dependent upon the AIMD increase parameter α .

IV. RELATED WORK

While numerous models of TCP have been developed for wired networks, few exist for wireless networks. Existing



(a) $\beta=0.5$ (b) $\beta=0.8$

Fig. 8. Convergence rate of TCP download flows with a wireless bottleneck (NS simulation, two TCP downloads, topology as in Figure 5).

802.11 wireless models are largely confined to static MAC layer properties such as transmission rate and collision probability. [5] develops a p-persistent model and uses this to study the quasi-polling behaviour of TCP downloads in an 802.11b WLAN while [9] develops a model of 802.11e MAC operation for TCP uploads. To the authors knowledge, a mathematical model of TCP dynamics, and analytic results on the convergence rate and RTT unfairness of mixed TCP upload and download flows in wireless networks, have previously been unavailable.

V. CONCLUSIONS

In this paper we develop an interaction model of TCP flows in a network with a wireless hop and illustrate its application in the context of 802.11e WLANs. Consideration encompasses characteristics of TCP flows such as RTT unfairness and responsiveness. We observe that TCP upload flows may exhibit quite different properties from TCP download flows. For example, RTT unfairness is absent in TCP uploads over wireless networks and convergence rates are insensitive to the AIMD backoff parameter in TCP.

VI. ACKNOWLEDGEMENTS

This work was supported by Science Foundation Ireland grant 03/IN3/I396.



(a) $\alpha_2 = 1, \beta_2=0.5$ (b) $\alpha_2=0.2, \beta_2=0.95$

Fig. 9. Convergence of TCP congestion windows of two TCP upload flows as AIMD α and β parameters of second flow are varied (NS simulations, topology is as in Figure 5 with $B=10\text{Mbs}$ yielding wireless bottleneck).

REFERENCES

- [1] H. Balakrishnan, V. Padmanabhan, "How Network Asymmetry Affects TCP". IEEE Communications Magazine, April 2001, pp60-67.
- [2] A. Detti, E. Graziosi, V. Minichiello, S. Salsano and V. Sangregorio, "TCP fairness issues in IEEE 802.11 based access networks", submitted paper.
- [3] S. Pilosof, R. Ramjee, Y. Shavitt, P. Sinha, "Understanding TCP fairness over Wireless LAN", INFOCOM 2003, 1-3 April 2003, San Francisco, USA.
- [4] H. Wu, Y. Peng, K. Long, S. Cheng, J. Ma, "Performance of Reliable Transport Protocol over IEEE 802.11 Wireless LAN: Analysis and Enhancement", INFOCOM 2002, 23-27 June 2002, New York, USA.
- [5] R. Bruno, M. Conti, E. Gregori, "Throughput Analysis of TCP Clients in Wi-Fi Hot Spot Networks", WONS 2004, 21-23 January, Trento, Italy.
- [6] G. Bianchi, "Performance analysis of IEEE 802.11 distributed coordination function", *IEEE Journal on Selected Areas in Communications*, 18(3):535-547, March 2000.
- [7] R. Shorten, F. Wirth, D. Leith, "A Positive Systems Model of TCP-like Congestion Control: Asymptotic Results". Hamilton Institute Technical Report 2004-1, April 2004.
- [8] R. Shorten, D. Leith, "Analysis and Design of Congestion Control in Synchronised Networks". Automatica, 41:725-730, April 2005.
- [9] D. Leith, P. Clifford, "Using the 802.11e EDCF to Achieve TCP Upload Fairness over WLAN Links", WiOpt 2005, Trento, Italy.
- [10] D. Leith, P. Clifford, "TCP Fairness in 802.11e WLANs", Wirelesscom 2005, submitted.
- [11] Berman, A., Shorten, R., Leith, D., "Positive Matrices Associated with Synchronised Communication Networks". Linear Algebra and its Applications, 393:47-54, 2004.
- [12] Padhye, J., Firoiu, V., Towsley, D., and Kurose, J., "Modeling TCP Throughput: a Simple Model and its Empirical Validation", Proc. SIGCOMM 1998.