

A Queueing Network Model of Short-Lived TCP Flows with Mixed Wired and Wireless Access Links*

Roberta Fracchia, Michele Garetto, and Renato Lo Cigno**

Dipartimento di Elettronica – Politecnico di Torino
Corso Duca degli Abruzzi, 24 – I-10129 Torino, Italy
{fracchia, garetto, locigno}@mail.tlc.polito.it

Abstract. We present an analytical model, based on a Fixed Point Approximation (FPA) solution, that can be used to derive the performance of different sets of TCP connections that share, and compete for, a common resource, typically a link and its associated buffer. A set of TCP connections is a group of connections that can be considered homogeneous, e.g., they have similar RTTs and all have a wireless access. TCP connections are modeled through the OMQN (Open Multi-class Queueing Network) paradigm. The conditions that define the feasibility of the solution and allow the convergence of the model are discussed and an application example with a RED buffer where wired and wireless connections converge is presented.

1 Introduction

After the explosion of cellular telephony, wireless and mobile networks are becoming more and more popular also for data transmission. Specialized network architectures, like WAP (Wireless Application Protocol) [1,2,3,4] and i-mode [5,6], have been proposed and deployed trying to overcome the inevitable impairments of wireless channels. These new architectures may finally find a position in the telecommunication market, but they do not seem able to substitute the traditional TCP/IP architecture for data transmission. The fast diffusion of wireless LANs and the standardization of new ultra-fast radio interfaces [7,8,9], together with the introduction of 3G cellular systems, is a strong indication that the use of standard, PC-based terminals running TCP/IP and accessing the Internet through a wireless connection will be a common scenario in a very short time.

The interest in the performance of TCP over wireless channels is thus on the rise, and the issue of how well wireless TCP connections will be able to compete with wired ones to share networks resources is matter of debate. Several simulation-based studies on wireless TCP have appeared in recent years [10,11], and some effort was also placed in trying to model the TCP behavior on wireless channels, capturing the essential feature

* This work was supported by the Italian Ministry for Education and Scientific Research through the PLANET-IP Project and by CERCOM, the Center for Multimedia Radio Communications of Politecnico di Torino

** Starting Nov. 1 2002 R. Lo Cigno is with the Department of Computer Science and Telecommunications (DIT) of the University of Trento, Italy

of random losses [12]. TCP in fact can't discriminate between packets lost due to lossy links, and those dropped due to congestion in the network.

The aim of this work is the extension of the modeling technique based on OMQN (Open Multiclass Queuing Network) models of TCP, combined with a Fixed Point Approximation (FPA) solution, to cover cases when different TCP connections experience different network conditions before competing for resources on one or more bottlenecks. One of the most interesting case is the wireless/mobile access, that can significantly modify the loss probability perceived by TCP connections, leading to unfair sharing of resources. Though this phenomenon is not new, to the best of our knowledge there are no other analytical approaches that allows to predict the performance and unfairness index in this situation.

2 Queuing Networks Models of TCP

Models of TCP in recent literature are numerous, with methods as different as heuristic analysis with empirical validation as in [13], or sophisticated mathematical abstractions as in [14]. A quick discussion of different approaches for TCP modeling can be found in [15] and its repetition here is superfluous. We have successfully used queuing models of end-to-end protocols, TCP being a particular case, in several works [15,16,17,18,19], so that the basic modeling method we adopt in this work is already extensively described, and we only recall here its main features.

An OMQN model is a queuing network model in which all queues are $M/G/\infty$. The customers of the OMQN represent TCP connections. The state of a TCP connection (customer) is uniquely identified by a pair (q, c) , where q represents a specific state of the protocol, described with a queue, and c identifies the number of remaining packets to be sent before the completion of the connection. An OMQN model is capable of describing any protocol whose dynamics can be described with a Finite State Machine (FSM). The use of classes to describe the backlog of the connection allows modeling short lived connections, which is indeed one of the main innovative characteristics of OMQNs. Given the average RTT of connections and the average packet loss probability, the OMQN model defines the load offered to the IP network by the aggregate of the connections represented by all the customers currently in the queuing network. New TCP connections enter the system with rate λ_{ext} and exit when reach class 0, i.e., they have completed the transmission. The arrival rate, together with the initial class of customers, that define the amount of information to be transferred, represent the nominal load of the system.

An OMQN model must be complemented with a suitable description of the network that TCP connections exploit to transfer packets. The network model must be capable, given the load, to compute the basic parameters that drive TCP performance: the average loss probability and the average RTT. The overall solution is obtained iterating the solution of the two sub-models with a Fixed Point Approximation (FPA) technique (see Fig. 1). When convergence is reached, the OMQN allows the computation of the throughput and the duration of connections.

Previous works on OMQNs concentrated mainly on the TCP model, somewhat disregarding the underlying IP network. As discussed in Sect. 3, this works concentrates

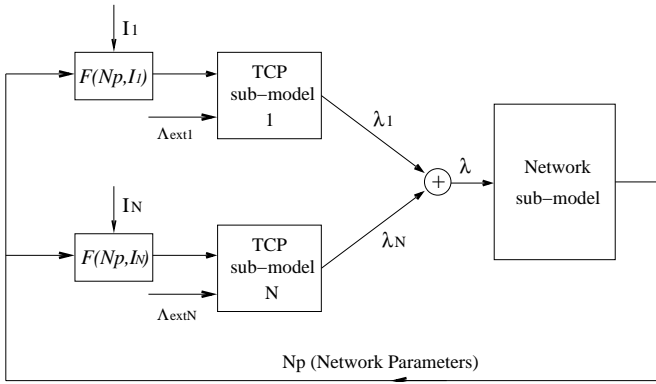


Fig. 1. Schematic representation of the Fixed Point Approximation method for the model decomposition, when N sets of TCP connections load the network

instead on the interaction of the OMQN models with the network, exploring the possibility of analysing the performance of TCP connections that experience different network conditions, but still have to compete for resources on a bottleneck. Sect. 4 presents a simple application of the method where two sets of TCP connections, one with a wired access and the other with a wireless access, compete for network resources on a RED-managed buffer. Results are validated against *ns-2* [20] based simulations.

We refer the reader to [15,16,17,18,19] for any further detail on the use of queueing networks in performance evaluation modeling of TCP connections.

3 Merging Different OMQN Models on a Bottleneck

One of the strengths of FPA techniques, is allowing the decomposition of an extremely complex modeling problem into simpler and partially decoupled sub-models, that interacts only through the iteration procedure. This is a classic application of the *divide at impera* (divide and conquer) principle. For instance the model decomposition allows using different modeling technique to tackle different tasks in the modeling process. Indeed, parts of the system could also be simulated or realized with a testbed instead of being analytically modeled.

Aiming at protocol performance evaluation, the natural problem decomposition is describing the protocol with one sub-model and the transport network with a different one. In the TCP/IP case, TCP becomes the protocol under analysis and the IP network, with everything that goes with it, the information transfer infrastructure. We stress the fact that the network sub-model is not a model of the IP protocol, but of the way the IP network reacts to the load offered by the TCP protocol.

An OMQN model allows the description of an *homogeneous* set of TCP connections, i.e., an ensemble of connections that share a similar propagation delay, the same maximum window size, the same version of the protocol (e.g., Tahoe, NewReno or SACK), etc. Here we use the OMQN TCP-NewReno model presented in [15].




Fig. 2. Network topology with a RED ($G1$) router and wired (D_i) and wireless (W_i) destinations

Let $\mathcal{N}()$ be the model of the network and $\mathcal{T}()$ be a generic OMQN model of a set of TCP connections. Given a synthetic workload $\bar{\lambda}$, a network model returns packet dropping probabilities, and packet transfer delays, in general some network parameters \overline{N}_p , while a TCP model returns a synthetic workload $\bar{\lambda}$ given \overline{N}_p . If TCP connections loading the network are not homogeneous, then $\bar{\lambda}$ and \overline{N}_p are, in the general case, all vectors.

The conditions for an FPA approach to converge is that there exist a single combination of $\bar{\lambda}$ and \overline{N}_p solving the system

$$\begin{cases} \overline{N}_p = \mathcal{N}(\bar{\lambda}) \\ \lambda_i = \mathcal{T}_i(N_p^i) \end{cases} \tag{1}$$

Fig. 1 represents a case where the workloads offered by TCP connection sets to the network simply sum one another and N_p^i are obtained from average values yielded by the network model \mathcal{N} through a modifying function $F(N_p, I_i)$ that includes some additional interference I_i , that does not enter the FPA convergence process. In every TCP sub-model there are independent external arrivals λ_{ext_i} ; $i = 1, \dots, N$.

In the following we shall use this general framework to model TCP connections mixed on the same bottleneck but with etherogeneous access links. Every TCP model is used to describe the behavior of connections following routes with different characteristics, represented by the functions $F(N_p, I_i)$.

4 Sample Applications

4.1 Network Topology

Consider a single bottleneck with a generic AQM (Active Queue Management) system, say RED to focus on a specific case, with TCP connections converging from different subnetworks: for instance a wired network through router **R** and a wireless LAN through a base station **BS**. Fig. 2 represents this simple, yet interesting scenario. Wired

connections can lose packets only in the bottleneck, while over the wireless channel there are further losses.

In these examples we assume the same round-trip-time for wired and wireless paths, in order to simplify the interpretation of results.

The network sub-model is described by a $M^{(X)}/M^{(X)}/1/B$ queue and represents an output interface of an IP router with a buffer of capacity B packets. Batch arrivals and batch services are introduced to model traffic burstiness. This representation was introduced in [15] for DropTail routers and it is modified here to manage RED AQM. It receives as input the traffic generated by TCP connections, given by the sum of the average loads offered by the different TCP sub-models. We suppose that the arrivals from each model are Poisson processes with rate λ_i with $i = 1, \dots, N$.

To model the clustering of transmissions due to TCP behavior we use batch arrivals. We assume that batches of packets arrive at the router with distribution $G = \{g_j\}$ where g_j is the probability that the size of a group is equal to j . The load generated by every model is:

$$\lambda_i = \sum_{j=1} j g_{j_i} \quad i = 1, \dots, N \tag{2}$$

where g_{j_i} is the dimension of the groups for each i TCP model. The average load λ offered by TCP is the sum of the number of the groups with the same dimension, in order to keep the traffic burstiness. It results:

$$\lambda = \sum_{i=1}^N \lambda_i = \sum_{i=1}^N \sum_{j=1} j g_{j_i} \tag{3}$$

So we have batch arrivals at the network queue and the probability gn_j that the size of a group is equal to j is given by:

$$gn_j = \sum_i g_{j_i} \tag{4}$$

TCP burstiness should be represented with an arrival correlation function within an RTT, and not simply by bursts arrivals, since packets never arrives at the same instant at the router buffer. Batch arrivals are a convenient mathematical approximation, that, however, ends up in a very pessimistic assumption. In order to compensate this pessimistic assumption, we use batch services. In [15] and [19] there are further explanations and validation tests for this assumption.

The IP network sub-model estimates the packet loss probability and the average round-trip-time, comprising queuing delay at the router as well as propagation delays. These parameters are fed back to the TCP sub-model in an iterative procedure that is stopped only when convergence is reached. We study the convergence of both loss probability calculated by the IP sub-model and loads offered to the network. As can be seen in Fig. 1, the input parameters of TCP sub-models are functions of the results of the network model (the same for each model) and of the scenario parameters I_i with $i = 1, \dots, N$, that may be different for every model. For these reasons if in two following steps the network parameters are the same, so that convergence is reached, and if the input scenario parameters I_i don't change, also the loads λ_i are the same. So when the

loss probability converges, also the different loads of the different TCP models reach the convergence since they are deterministic functions of the network parameters. It's important to stress that every analyzed parameter reach the convergence at the same time even if the whole model is asymmetric: at the entrance of the IP network model there is a sum of the results of different TCP models while the results of IP sub-model are the input parameters of each TCP sub-model.

Thanks to the use of different TCP sub-models, we can run the OMQN to calculate the average loss probability and the completion time of wired as well wireless connections, even if they are mixed.

4.2 Network and TCP Parameters

In this single-bottleneck topology we use a 45 Mb/s link whose length is equal to 5,000 km. We assume one-way TCP connections with uncongested backward path, so that ACKs are never lost or delayed at queues. The distance of servers from the congested router is assumed to be uniformly distributed between 200 and 3,800 km, while the distance of the wired clients on the other side of the link is assumed to be uniformly distributed between 200 and 2,800 km. The wireless clients distance from BS is negligible. Connections are established choosing at random a server-client pair, and are generated following a Poisson process. The packet size is constant, equal to 1,000 bytes; the maximum window size is 32 packets. The TCP tic is equal to 500 ms. The buffer of the $M^{(X)}/M^{(X)}/1/B$ queue is a RED buffer to avoid global synchronization and biases against bursty traffic. The size of the buffer is 128 packets, we use the 'gentle' algorithm with a minimum threshold of 10 packets, a maximum threshold of 50 packets and a queue weight $w_q = 10^{-5}$.

4.3 Bernoulli Loss Process

In this case, the loss probability P_{Lw} of the wireless link is considered Bernoulli. The average loss probability for wireless connections is:

$$P_W = 1 - (1 - P_L)(1 - P_{Lw}) = P_L + P_{Lw} - P_L P_{Lw} \quad (5)$$

where P_L is the loss probability estimated by the network model and due to the bottleneck. In other words the further interference I_i for wireless connections consists of an additional packet loss probability only, and the function $F()$ has the form expressed in (5). We show the behavior of the protocol for different rates of wireless links and for different number of wireless connections. The random loss probability on wireless channels is $P_R = 10^{-2}$.

Numerical Results Performance figures are plotted versus the total normalized nominal load (link utilization that would result without any retransmissions), for transfers of 100 segments objects. Results are validated against point estimates and 95% confidence intervals obtained with $ns - 2$ simulations [20] with a confidence level of 90%. Figs. 3 and 4 show the loss probability and the average completion time of both wired

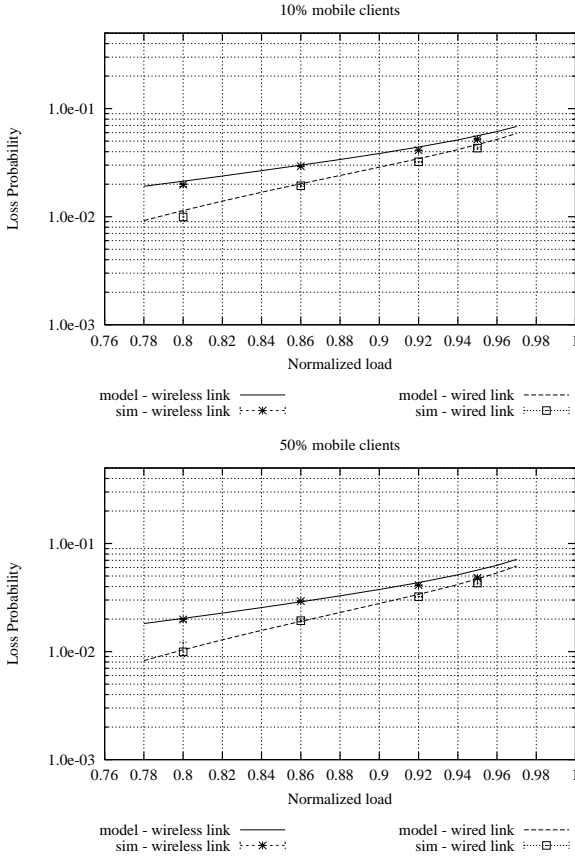


Fig. 3. Loss probability; Bernoulli wireless link: 45 Mbit/s, $P_R = 10^{-2}$

and wireless connections when the rate of wireless link is 45 Mbit/s, like in the recently standardized 802.1q VLAN (Virtual Local Area Network). Upper plots refer to a case when wireless clients are 10% of the total number while in bottom plots they are 50%. The model give accurate results differentiating wireless and wired connections, following closely values obtained with *ns-2* simulations. Fig. 5 and Fig. 6 presents the same figures when the rate of wireless links is 11 Mbit/s, like in standard 802.11b LANs.

We notice incidentally that RED AQM is not able to enforce any degree of fairness between wired and wireless connections. RED droppings should be proportional to the route of the TCP flows, thus balancing the performance among connections. With finite flows, however, the number of packets to be transferred remains constant and RED is not able to differentiate between wireless, low throughput users, and wired ones, that have a throughput 40-50% higher. This result may sound obvious; however it was never stressed in RED literature and the OMQN/FPA modeling approach offer a theoretical explanation of the phenomenon. In fact it is easy to demonstrate through it that a well-

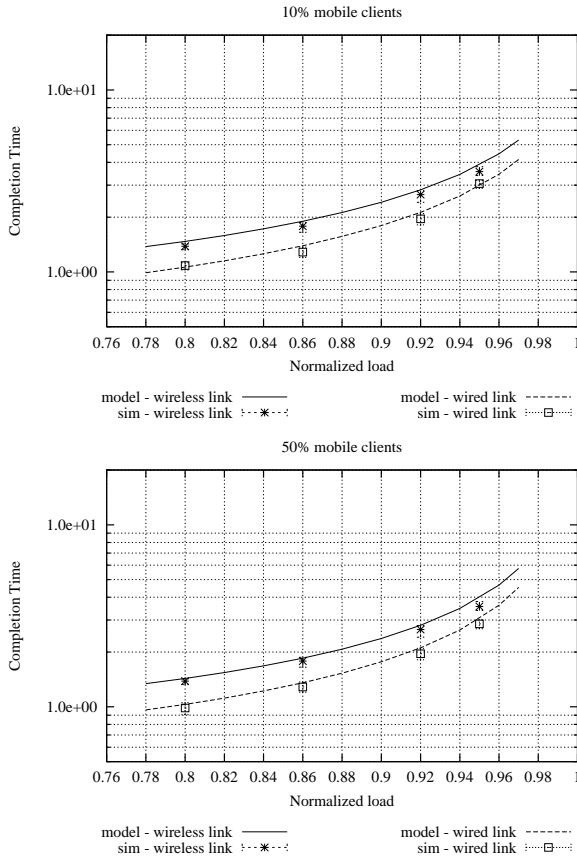


Fig. 4. Average Completion time; Bernoulli wireless link: 45 Mbit/s, $P_R = 10^{-2}$

tuned RED buffer (i.e., one where forced losses never occur) introduces losses that are uncorrelated, hence introducing a Bernoulli loss process. This implies dropping on average an equal number of packets from any connections with the same length, regardless of the actual throughput. Flow lengths of 20, 50 and 200 packets yield similar results, not reported to avoid cluttering the figures.

4.4 Two State Gilbert Wireless Channel

Over many real channels the errors occur in bursts, separated by fairly long error free gaps. To represent this situation, we use the Gilbert channel model: a binary channel described by a two state Markov chain. As shown in Fig.7, we model it with a Continuous Time Markov Chain (CTMC), in which in the Good State (G) errors occur with probability P_G , while in the Bad State (B) they occur with higher probability P_B . The transition matrix Q is:

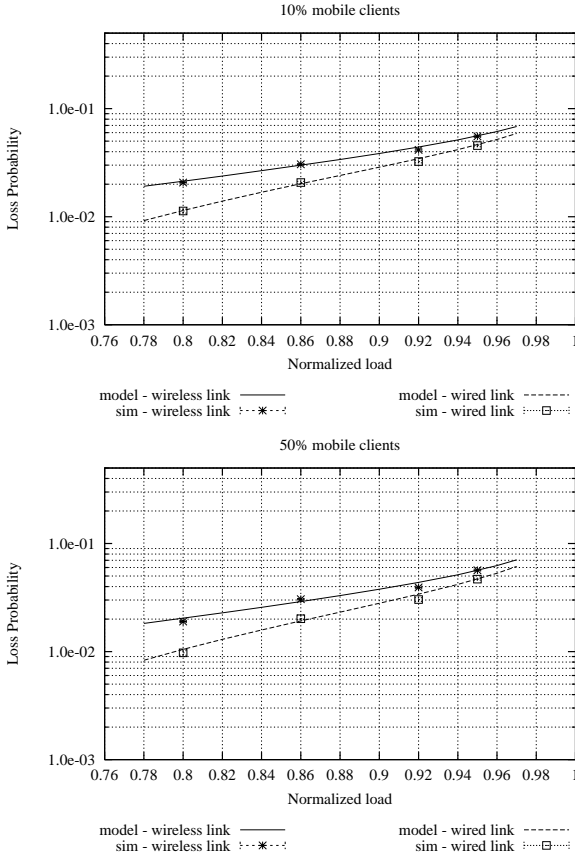


Fig. 5. Loss probability; Bernoulli wireless link: 11Mbit/s, $P_R = 10^{-2}$

$$Q = \begin{bmatrix} -\alpha & \alpha \\ \beta & -\beta \end{bmatrix}$$

where α is the rate of transition from G to B , while β from B to G . The probabilities of being in Good and Bad States are: $\pi_G = \frac{\alpha}{\alpha + \beta}$ $\pi_B = \frac{\beta}{\alpha + \beta}$ The sojourn times for states B and G are exponentially distributed with means: $E\{\tau_G\} = \frac{1}{\alpha}$ $E\{\tau_B\} = \frac{1}{\beta}$ We consider $P_G = 0.01$ and $P_B = 0.1$ and different probabilities of being in Good or Bad.

Referring to Fig. 1, a Gilbert channel requires the use of three different TCP models: one for wired connections, one for wireless connections in the bad state and one for the good state. The average loss probability P_{W_G} for flows being in Good or Bad channel state are computed using (5). Hence, the average loss probability for wireless connections is $P_W = P_{W_G}\pi_G + P_{W_B}\pi_B$




Fig. 6. Average Completion time; Bernoulli wireless link: 11Mbit/s, $P_R = 10^{-2}$




Fig. 7. Gilbert Model

Fig. 9 and Fig. 8 show the loss probability and the average completion time of both wired and wireless connections, 100 packets long, when the rate of wireless link is 11 Mbit/s and mobile clients are 50% of the total number. We consider two different cases: in the first, the probability of being in B is $\pi_B = 0.01$, in the second it is $\pi_B = 0.1$. Also

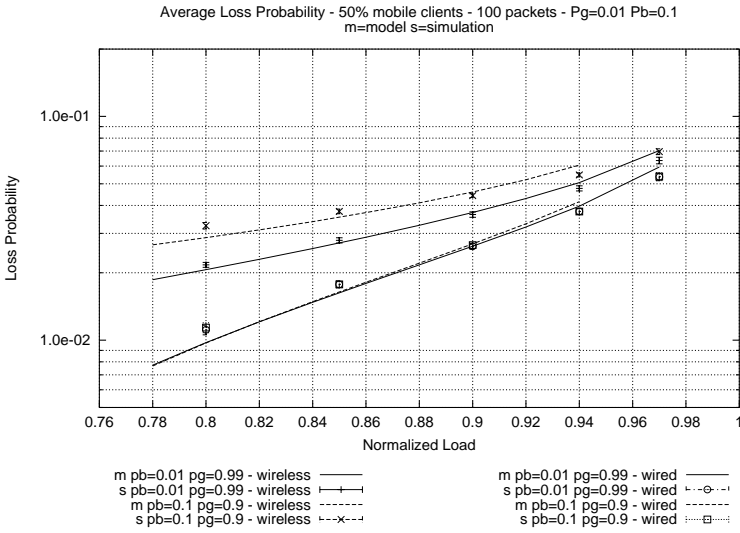


Fig. 8. Loss probability; Gilbert wireless link: 11Mbit/s

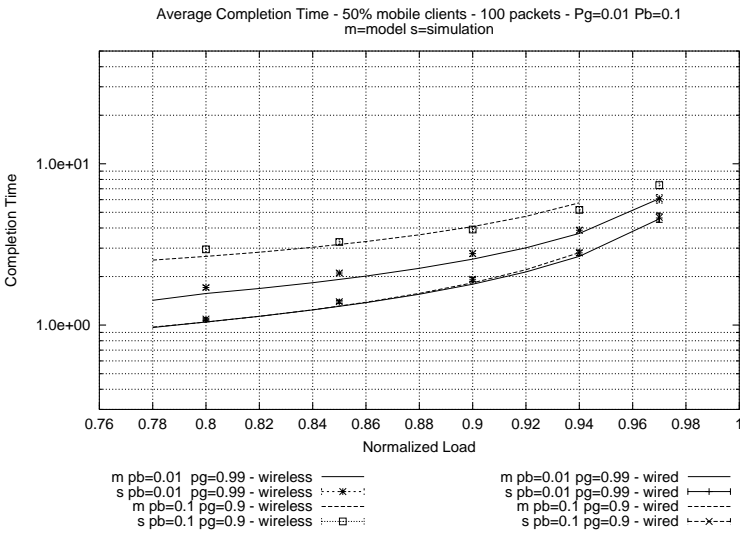


Fig. 9. Average Completion time; Gilbert wireless link: 11Mbit/s

in this more complex case, results provided by the model match closely the simulation points.

5 Conclusions and Future Work

This paper has presented a method to derive the performance of different sets of TCP connections when are competing for a common resource in the Internet. A *set* of TCP connections is defined as an ensemble of connections that share common characteristics, such as the average RTT, the maximum window size, etc. and each set of connections can be modeled through an independent OMQN model. Workloads offered to the network by different connections are summed together to obtain the performance of the IP network.

We have shown an application of the method with connections that compete for a RED buffer. Different OMQN models represent sets of connections that either access the network through a lossy, wireless link or through a traditional, reliable wired network. Results were validated against *ns-2* simulations, showing that the method is viable and yield accurate results. Results show that, as expected, wireless connections are heavily penalized at any load of the bottleneck. Simple AQM disciplines like RED, do not seem able to provide any protection against this penalization.

References

1. O. Kone, "Experiment with the Validation of WAP Systems," in *Proc. IEEE LCN'00*, Nov. 2000, Tampa, Florida
2. D. Maierhofer, "Internet and WAP Technology as Key for Easy Service Management and Convenient Call Control," in *Proc. Intelligent Network Workshop'00*, 2000
3. C.L.C. Wong, M.C. Lee, R.K.W. Chan, "GSM-Based Mobile Positioning Using WAP," in *Proc. Wireless Communications and Networking Conference'00*, Sept. 2000, Chicago
4. A. Andreadis, G. Benelli, G. Giambene, B. Marzucchi, "Performance Analysis of the WAP Protocol over GSM-SMS," in *Proc. IEEE ICC'01*, June 2001, Helsinki, Finland
5. G. Vincent, "Learning from i-mode [Packet-based Mobile Network]," *IEEE Review*, Vol. 47, Issue 6, Nov. 2001
6. K. Enoki, "i-mode: the Mobile Internet Service of 21st Century," in *Proc. Solid-State Circuits Conference'01*, Feb. 2001, San Francisco, CA
7. M. Natkaniec, "Wireless LAN Standards and Applications," *IEEE Communications Magazine*, Vol. 40, Issue 2, Feb. 2000
8. T. Ichikawa, "IEEE802: Compliant High-Speed Wireless LAN", in *Proc. High Performance Switching and Routing'02*, May 2002, Kobe, Japan
9. Haitao Wu, Yong Peng, Keping Long, Shiduan Cheng, Jian Ma, "Performance of Reliable Transport Protocol over IEEE 802.11 Wireless LAN: Analysis and Enhancement," in *Proc. IEEE INFOCOM'02*, June 2002, New York
10. A. Chockalingam, M. Zorzi, V. Tralli, "Wireless TCP Performance with Link Layer FEC/ARQ," in *Proc. IEEE ICC'99*, June 1999, Vancouver
11. G. Racherla, S. Radhakrishnan, C.N. Sekharan, "Performance Evaluation of Wireless TCP Schemes under Different Rerouting Schemes in Mobile Networks," in *Proc. IEEE TENCON'98*, Dec. 1998, New Delhi
12. H. Balakrishnan, N. Padmanabhan, S. Seshan, R.H. Kats, "A Comparison of Mechanisms for Improving TCP Performance over Wireless Links", in *Proc. SIGCOM'96*, Oct. 1996, Stanford, CA, USA
13. J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP Throughput: A Simple Model and its Empirical Validation," in *Proc. ACM SIGCOMM'98*, Sept. 1998, Vancouver

14. M. Grosslauser, J. Bolot, "On the Relevance of Long-Range Dependencies in Network Traffic", *IEEE/ACM Transaction on Networking*, Vol.7, N.5, pp.629–640, Oct. 1999
15. M. Garetto, R. Lo Cigno, M. Meo, E. Alessio, M. Ajmone Marsan, "Modeling Short-Lived TCP Connections with Open Multiclass Queueing Networks," *7th International Workshop on Protocols For High-Speed Networks (PfHSN 2002)*, Berlin, Germany, April 22-24, 2002.
16. R. Lo Cigno, M. Gerla, "Modeling Window Based Congestion Control Protocols with Many Flows", *Performance Evaluation*, No. 36–37, pp. 289–306, Elsevier Science, Aug. 1999
17. M. Garetto, R. Lo Cigno, M. Meo, M. Ajmone Marsan, "A Detailed and Accurate Closed Queueing Network Model of Many Interacting TCP Flows," in *Proc. IEEE Infocom'01*, April 22–26, Anchorage, Alaska, USA
18. E. Alessio, M. Garetto, R. Lo Cigno, M. Meo, M. Ajmone Marsan, "Analytical Estimation of the Completion Time of Mixed NewReno and Tahoe TCP Traffic over Single and Multiple Bottleneck Networks," in *Proc. IEEE Globecom 2001*, San Antonio, TX, USA, Nov. 25–29, 2001
19. M. Garetto, R. Lo Cigno, M. Meo, M. Ajmone Marsan, "On the Use of Queueing Network Models to Predict the Performance of TCP Connections," *Proc. 2001 Tyrrhenian International Workshop on Digital Communications*, Taormina (CT), Italy, Sept. 17–20, 2001
20. ns-2, *Network Simulator (version 2)*, Lawrence Berkeley Laboratory, <http://www.isi.edu/nsnam/ns>.
21. T. Bu, D. Towsley, "Fixed Point Approximation for TCP Behavior in an AQM Network," in *Proc. ACM SIGMETRICS 2001*, June 16-20, Cambridge, MA, USA
22. W. R. Stevens. *TCP/IP Illustrated, vol. 1*. Addison Wesley, Reading, MA, USA, 1994.