periodica polytechnica

Electrical Engineering
51/1-2 (2007) 65–71
doi: 10.3311/pp.ee.2007-1-2.07
web: http://www.pp.bme.hu/ee

RESEARCH ARTICLE

Enhanced congestion control in TCP for solving hidden terminal problems in ad hoc wireless networks

Mohamed Yahia / *József* Bíró

Received 2007-06-05

Abstract

This paper studies TCP performance over multihop wireless ad hoc networks that use the IEEE 802.11 *protocol as the access method. The aim is to improve the TCP fairness while keeping the algorithm as simple as possible, since in previous works the algorithm designs were more complicated. We propose a simple approach to improve fairness based on scheduling (pacing) new packets according to the transmission interval formed from scaled round-trip time (RTT) and congestion window. Our simulation shows that, given specific scale parameter x, TCP achieves high fairness and throughput via improved spatial channel reuse, if it operates in a certain range of the transmission interval.*

Keywords

ad hoc networks · wireless networks · TCP protocol

Mohamed Yahia

Department of Telecommunications and Media Informatics, BME, H-1521 Budapest, Hungary

József Bíró

Department of Telecommunications and Media Informatics, BME, H-1521 Budapest, Hungary

e-mail: biro@tmit.bme.hu

1 Introduction

Wireless networks are rapidly becoming common place; upcoming technologies support wireless communication on chipsets; new laptops have built-in WiFi cards; hotspots are being installed in airports, hotels and coffee shops; and offices are converting their existing local area networks to wireless to allow their workers take advantage of mobility. In packet-radio networks a medium access control (MAC) protocol is essential so that stations can share a common broadcast channel. Carrier Sense Multiple Access (CSMA) protocols have been used in a number of packet-radio networks. The goal of these protocols is to prevent multiple stations from transmitting simultaneously within their range, by listening on the channel before transmitting. CSMA can not solve hidden-terminal problems which cause degradation in the performance, however, because it cannot prevent collisions. A hidden terminal is a node in the receiver's neighbourhood which *cannot* detect the sender and may disrupt the current packet transmission [2]. Many solutions have been suggested to solve this problem at the level of routing and Transmission Control Protocol (TCP) protocols [5]. In this paper we introduce a solution to the hidden terminal problem on the level of the TCP protocol. M-TCP (Modified TCP) is a novel and simple congestion control algorithm with TCP over multihop IEEE 802.11 networks. It implements rate-based scheduling of transmission within the TCP congestion window. Rather than adaptively estimating the 4-hop propagation delay and the coefficient of variation of recently measured round-trip times as in [4], in our approach a TCP sender sets its transmission interval for the current congestion window and round-trip time by using a scaling parameter x. The useful range of x can be identified in a straightforward manner. A comprehensive study using the NS2 Network Simulator[1] shows that M-TCP achieves a fairness greater than 99%, provides high throughput in almost all scenarios, and is highly responsive to changing network traffic. Our approach provides each node with a fair share of the available bandwidth, even when flows are not within each other's transmission range but are within each other's interference range. M-TCP is therefore a good solution to the hidden

[1] http:/www.isi.edu/nsnam/ns/doc/node218.html

terminal problem.

We compare our algorithm to other TCP algorithms for wireless networks such as TCP Westwood, Adaptive TCP (TCP-AP), and also to standard TCP (TCP Reno). In order to isolate the IEEE 802.11 MAC-induced deficiencies of TCP in multihop wireless networks, we only consider static scenarios in this paper.

This paper is organized as follows. Background information and review previous work is provided in Section 2. In Section 3 we present the M-TCP congestion control algorithm. We describe the simulation environment and results in Section 4. Section 5 compares the simulated performance of M-TCP to that of other TCP variants. We conclude the paper in Section 6.

2 Background and Related Work

TCP is a connection-oriented transport layer protocol that provides reliable in order delivery of data to a TCP receiver. Since the characteristics of wire-line and wireless networks differ, a TCP algorithm which is designed to perform well in wireline networks suffers from degradation in wireless networks. One of the problems with TCP in ad-hoc networks is the effect of a high bit error rate, which may corrupt packets and result in lost data segments or acknowledgments (ACKs). If an acknowledgment is not received by the sender within the retransmission time out (RTO) window, the sender does the following:

1 Retransmits the segment.

2 Exponentially backs off its retransmission timer.

3 Reduces its congestion control window threshold.

4 Closes its congestion window to one segment.

In the case of repeated errors the congestion window at the sender will remain small, resulting in low throughput. When hidden terminals exist in the network and the destination node is located in their interference range, a sender node (B) that receives no response after seven retransmissions of the RTS will drop its head-of-line packet according to the IEEE 802.11 MAC layer. Waiting for a period of time longer than the RTO results in the TCP sender timing out, retransmitting a packet and invoking congestion control.

Routing problems are another issue in ad hoc networks, since they can cause TCP degradation. Taking into consideration the differences between mobile and stationary networks, one finds that routing problems rarely occur in static networks. Routing problems are thus not considered in this work.

TCP Westwood (TCPW) [6] is essentially designed to improve the performance of TCP Reno in both wired and wireless networks. TCPW relies on end-to-end bandwidth estimation to determine the cause of packet loss (congestion or wireless channel effect), which is a major problem in TCP Reno. The idea is to continuously evaluate the quality of the connection at the TCP source by monitoring the rate of returning ACKs, and set the congestion window and slow start threshold ($ssthresh$) accordingly. Although TCPW performs well in ad hoc wireless networks, it still suffers from unfairness when hidden terminals are present.

TCP Adaptive Pacing (TCP-AP) is another version of the TCP protocol which tries to solve the hidden terminal problem in ad hoc wireless networks. Sherif et al. [4] introduced TCP-AP, which implements rate-based scheduling within TCP's congestion window in order to avoid burst packet transmissions and quantify incipient congestion. It measures the fluctuation (coefficient of variation) in a sample of round trip times and adaptively calculates the appropriate pacing of transmission. TCP-AP schedules new packets according to this computed rate, and allows the node to delay transmission of a data packet until the previously sent packet has been forwarded 4 times. Using simulations the authors showed that TCP-AP could substantially improve goodput with respect to TCP NewReno.

Singh et al. [10] developed a MAC protocol that employs adaptive interference cancellation based on cross-layer design considerations to increase network throughput and fairness. They used simulations with different tools and conclude that in terms of providing throughput gains and energy savings, it is better to exploit multiuser diversity by interference cancellation than to use multiple antennas with an 802.11a MAC.

Chane et al. [3] analysed Floor Acquisition Multiple Access (FAMA) protocols for single-channel packet-radio networks with hidden terminals. These protocols permit a station to acquire control of the channel dynamically before transmitting data packets. Their verification and throughput analysis, also supported by simulations, demonstrated that carrier sensing significantly improves performance in single channel networks in the presence of hidden terminals.

As opposed to the approaches mentioned above, our enhanced TCP requires neither modifications to the routing or link layers, nor cross-layer information from intermediate nodes along the path. Our M-TCP approach also fundamentally differs from TCP-AP [4]. Although M-TCP does depend on the delay of transmission packets, the computation of transmission intervals is based on the measured congestion window and round-trip times rather than using 4-hop transmission delay and the RTT coefficient of variation.

3 Rate-Based Transmission of Packets

3.1 Motivation

In IEEE 802.11, control handshake Request-To-Send/Clear-To-Send(RTS/CTS) messages precede each packet transmission. Due to the spatial reuse constraint of the wireless channel, neighbouring nodes of both sender and receiver defer their transmission until the subsequent DATA-ACK transmission is completed. Thus packet bursts caused by TCP's window-based congestion control result in increased contention on the wireless channel. This contention on the link layer may lead to packet drops due to the hidden terminal problem [9]. Our TCP algo-

rithm focuses on improving fairness in the presence of a hidden terminal. Since the sender can obtain rate information from the feedback packets of the receiver, the packet transmission in these networks can in principle be adapted to the rate. Our goal is to incorporate a rate-based transmission algorithm into a window-based TCP congestion control.

To this end, for each received ACK, TCP computes the transmission interval according to the current congestion window and the measured round-trip time RTT divided by a scaling parameter x. It then schedules the transmission of new packets based on this interval. This creates a rate-based congestion control which is sufficiently simple for efficient implementation.

In current TCP variants such as Reno and NewReno, congestion is managed solely upon the observation of packet losses. Although Vegas uses both RTT and packet losses to identify congestion, it still suffers from the negative effect of packet transmission bursts in wireless multihop networks. These characteristics of IEEE 802.11 make it obvious that TCP suffers from poor performance. Thus, the aim of this work is to develop a congestion control that identifies high contention on the network path of the TCP connection and throttles the transmission rate before losses occur. In order to improve the fairness of TCP on the network path, we propose a transmission interval adjustment based on the state of the congestion window CW, the recently measured round-trip time RTT, and a scale parameter x. This transmission interval is given by the formula:

$$T_{\text{interval}} = RTT \times 1/x \times CW$$

3.2 The effect of the scale parameter x on TCP performance

The scale parameter x, which can be selected based on measurements of the network performance, plays a very important role in our algorithm. Higher values of x lead to smaller intervals, which in turn increase the congestion window, improve spatial reuse, and enhance throughput. The user can set x to obtain the required level of fairness and throughput. Larger values of x yield higher throughput, but degrade the fairness. In this paper we consider which values of x result in both a high level of fairness and acceptable throughput.

3.3 The hidden terminal problem

A hidden terminal is a potential sending node in the receiver's interference range which cannot detect the sender, and thus may disrupt transmission of the current packet. Although CSMA is designed to solve the collision problem in an IEEE 802.11 wireless network by RTS/CTS handshake messages, a hidden terminal degrades its performance substantially. Carrier sensing simply cannot prevent these collisions.

Let us assume that we have chain network topology as in Fig. 1, where node 1 wishes to transmit to node 2 and node 4 wishes to transmit to node 5. In this case, node 4 is a hidden terminal for node 2 because it can neither receive the RTS/CTS handshake between nodes 1 and 2 nor sense the transmission from node 1 to node 2. Since node 1 is out of the sensing range of node 4, it can transmit to node 2 while node 4 transmits to node 5. Thus, a collision occurs at receiving node 2, since it is in the interference range of node 4, which leads to contention loss. We can also examine the hidden terminal problem in a cross topology (Fig. 2). Assume that node 2 transmits to node 9, and node 6 transmits to node 4. In this case node 4 becomes a hidden terminal for the transmission from node 2 to node 9.

In our approach, the hidden terminal problem can be avoided by adjusting the transmission interval of packets (T_{interval}).

3.4 Unfairness due to Hidden Terminals

The Distributed Coordination Function (DCF) in IEEE 802.11 defines two methods in accessing the medium; the two-way and four-way handshake. The two-way handshake occurs when sender transmits the data to the receiver and the receiver responds with an ACK if the data received successfully. In the four-way handshake, the sender first transmits out the RTS and responding to this; the receiver sends out CTS if it found the idle medium, after that the sender sends the data and receiver responds with an ACK. Every node in IEEE 802.11 allowed to maintain a contention window (w) and a back-off timer according to the well-known Binary Exponential Back-off (BEB) algorithm used. Whenever node wants to transmit, it defers by back-off timer which generated according to this formula:

$$Backoff - Timer = Random() \times SlotTime$$

Where Random value is uniformly distributed over the range [0, w] the SlotTime is specified by the physical layer. The new back-off timer value is generated whenever it becomes zero. At the first transmission of the packet the w will be set to $w_{\min}$ and doubles whenever retransmission is initiated. When retry limit is reached or in case of successful transmission the w will be reset to $w_{\min}$.

Now let us see what happens when node (say 4 in Fig. 1) wants to transmit. In this case, all the nodes in its range will freeze their back-off timers until the transmission completes and the medium becomes an idle. Then, all the nodes defer for DCF Inter-Frame Space (DIFS) period. While the node 4 generates new random value and backs off before it initiates another transmission, the other nodes resume to count down from their frozen back-off timers. Thus, node 4 may transmit several packets before another node's (e.g. node 2 when it sends CTS to node 1) back-off timer is reduced to zero. This is unfair for node 1 and may lead to starvation if it is repeated several times [7].

4 Simulation Setup and Results

The results reported in this paper are obtained using the *NS2* network simulator.

4.1 Setup Environment

The simulation parameters are given in Table 2. We simulated two static (fixed) networks of i and k nodes in the linear chain and cross configurations respectively (Figs 1 and 2). For each network we consider the performance in the absence and the presence of a hidden terminal. The specific simulation scenarios are shown in Table 1. In order to test our approach, we established two TCP connections with 1 kB packets continuously transferred by File Transfer Protocol (FTP). In all network scenarios, each node is separated by 250 meters from its adjacent nodes. Our ad hoc routing protocol is On-Demand Distance Vector Routing (AODV) [8]. Each simulation ran for 200 s. All simulation parameters which are not explicitly stated in this paper are set to the *NS2* defaults.

Tab. 1. Simulation scenarios

No	Scenario	Net. Type
1	node 1 send to 2 and 5 to 6	chain
2	node 1 send to 2 and 4 to 5	chain
3	node 1 send to 5 and 6 to 9	cross
4	node 2 send to 9 and 6 to 4	cross

Tab. 2. Simulation Setup parameters

Parameter	Type
transmission range	250m
interference, sensing range	550m
propagation model	Two-Ray Ground
antenna model	Omni-directional antenna
channel bandwidth	2Mbit/s
buffer size	50 packets

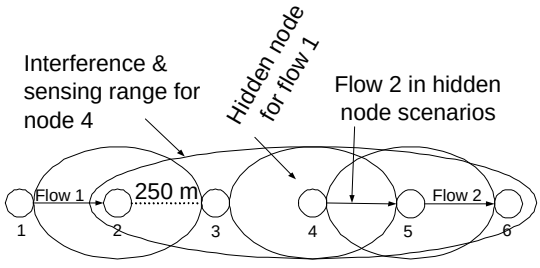

Fig. 1. Chain topology

4.2 Results

4.2.1 TCP Fairness

We begin by computing Jain's fairness index [1] for each parameter set:

$$F(bw_i, i = 1, .., n) = (\sum_{i=1}^{n} bw_i)^2 / n \sum_{i=1}^{n} bw_i^2,$$

where bw_i are the bandwidth shares of the TCP flows. Fig. 3 shows the impact of x on the fairness of the network.

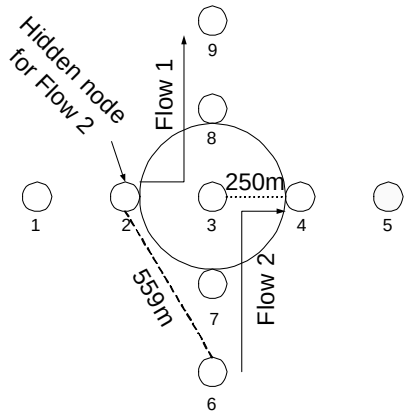

Fig. 2. Cross topology

Chain Network In Scenario 1 there is no contention due to a hidden terminal, so both flows take up equal shares of the available bandwidth. Thus, most TCP algorithms have a high fairness in this case. As shown in Fig. 3,[2] M-TCP achieves a fairness between 99% and 100% for all $x \leq 50$. The situation is different in Scenario 2, where TCP suffers from unfairness due to the interference caused by hidden terminals and the performance of the network degrades. M-TCP, however, continues to achieve a fairness index of more than 99% in this case for all values of x.

Cross Network In the cross network, flows share the forwarding node at the intersection. The fairness thus fluctuates around specific points. As shown in Fig. 3, the highest fairness levels were recorded for $x = 35$ in both Scenario 4 and Scenario 3 (98% and 97% respectively). Other values of x also achieve a fairness index of 95% or better.

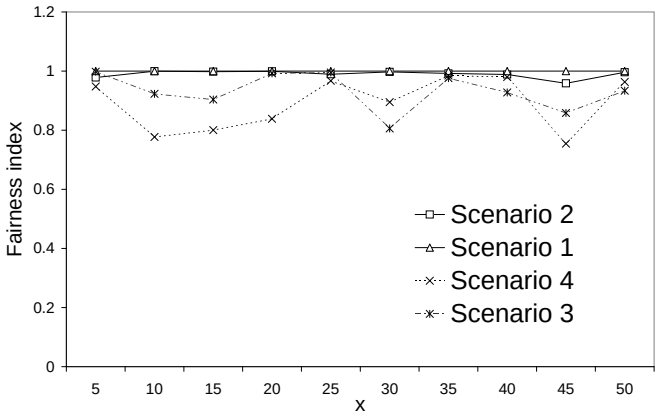

Fig. 3. M-TCP Fairness in all Scenarios

4.2.2 TCP Throughput

Although network fairness is the main interest of this paper, we also wish to investigate how well M-TCP achieves spatial channel reuse (since improving spatial reuse increases TCP throughput). In Figs. 4 and 5 we show both TCP flows in order

[2] In the figures, each tick mark means different simulation. We connect them to see the shape of the graph.

to observe the fairness. The throughput is shown as a function of the scaling parameter x for the chain and cross networks respectively. We observe that M-TCP flows share the available bandwidth equally in networks containing a hidden terminal, rather than allowing one flow to dominate as in other TCP variants (TCPW, TCP-AP, Standard TCP; see Figs. 8 to 11).

Chain Network In scenario 1, the flows are identical and the simulation shows that throughput increases as long as x increased (see Fig. 4) where it reaches the other TCP variants as it will be shown in Section 5. In contrast, the flows in Scenario 2 are "twisted". At $x = 50$, however, M-TCP throughput exceeds some of TCP versions by 1.5% or more (see Fig. 12 in Section 5).

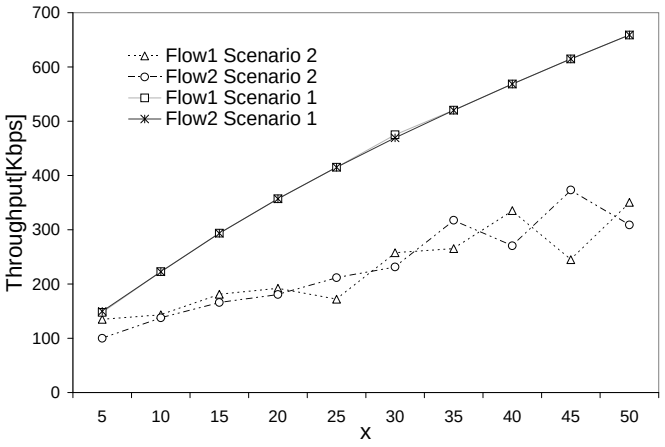

Fig. 4. Throughput in chain network

Cross Network In order to test the influence of a hidden terminal in the cross network we investigate Scenario 4 where node 4 (receiver side of flow 2) is under the effect of interferences of node 2. In Fig. 5, M-TCP improves the throughput flow 2 and achieves high fairness only for certain values of the scaling parameter ($x = 25$, 35, 40, and 50); note that these values are also evident in Fig. 3. In Scenario 3, both flows increase their throughput gradually with increasing x but in a "twisted" manner. When $x = 50$, the M-TCP throughput improves to 161 Kbps. This is 5% more than the value achieved by TCPW and standard TCP (152 Kbps). The maximum fairness achieve by M-TCP is also 97%, better than the other TCP variants (see Table 3).

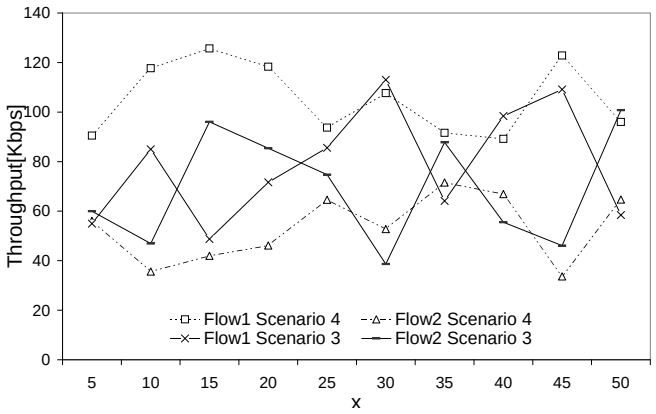

Fig. 5. Throughput in cross network

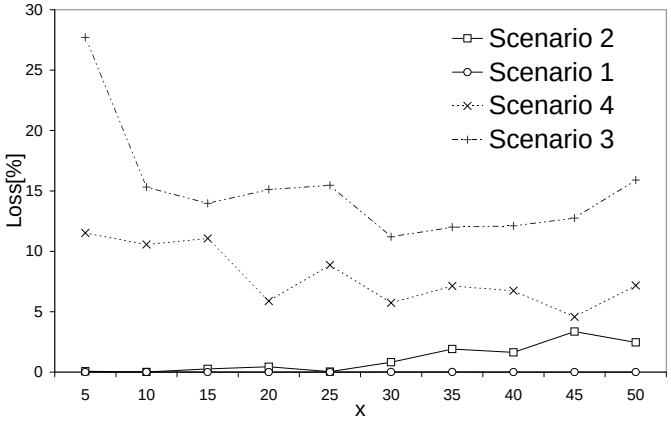

Fig. 6. Loss in chain and cross networks

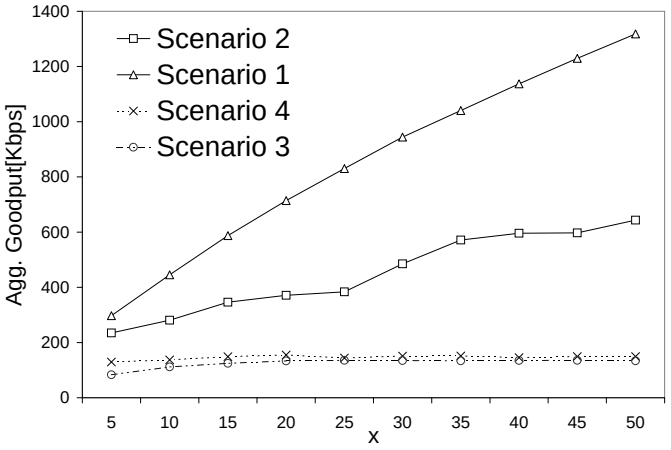

Fig. 7. Goodput in chain and cross networks

4.2.3 TCP Goodput and Loss

In a hardwired Internet, packet loss is mainly due to buffer overflow at the bottleneck router. In a multihop wireless network, on the other hand, packet drops may be caused by both buffer overflow and contention due to hidden terminals. In Figs. 6 and 7, we present the goodput and loss of all scenarios for those values of x where the network achieved the highest fairness index. We can observe an approximately linear increase of goodput with increasing x in scenarios 1 and 2 (Fig. 7). The loss in both of these scenarios (Fig. 6) is also quite small; even in Scenario 2, its maximum value is 3.35% at $x = 45$.

The goodput levels in scenarios 3 and 4 are very close to each other, and do not appear to depend strongly on x. Choosing the right value of x can still minimize the loss, however. In scenario 4, for example, choosing $x = 45$ results in a low loss rate of 4.5%.

For further improvements we have also undertaken a preliminary investigation of dropped packets. We find that in all scenarios most of the dropped packets are due to MAC layer collisions. As an example, consider the cross network without hid-

den terminals (scenario 3) at $x = 5$, where the highest rate of dropped packets occurred (Fig. 6). In this case we find that 12% of lost packets were dropped due to the callback in routing protocol, 77% were due to collisions in the MAC protocol, 6% were caused by route failure, 4% were due to retransmission time out, and 1% were dropped because of buffer size. Even this loss, however, can be reduced to lower rates by choosing larger x but on the expense of fairness.

5 Comparison with other TCP versions

We now compare the fairness obtained by M-TCP, TCP Westwood, TCP-AP, and Standard TCP for all four scenarios (Table 3). Note that M-TCP is always the fairest algorithm, sometimes by a substantial amount. In Figs. 8 through 11, we compare the behaviour of the congestion window over time under M-TCP, TCP Westwood, and TCP-AP and omit which is related to Standard TCP because it is identical to TCP Westwood. Since we are mainly concerned with hidden terminal problems, only the results obtained in Scenarios 2 and 4 are shown.

In Scenario 2 (Fig. 8), it is clear that under TCP Westwood flow 1 never had the opportunity to transmit at the maximum negotiated rate, because its congestion window was always significantly smaller than the advertized window size. The other flow thus dominated the channel bandwidth in this case. Under TCP-AP (Fig. 10), flow 1 made some attempts at transmission but never achieved a significant bandwidth share.

In Scenario 4 (Fig. 9), it is obvious that both TCP Westwood flows had better opportunity to transmit and this is also clear from the high fairness shown in Table 3. In TCP-AP (Fig. 11) flow 2 does not share the available bandwidth with flow 1 and results in low fairness. In Fig. 12 we show the aggregate throughput for all TCP versions. It is clear that M-TCP at $x = 50$ can reach the other TCP throughput or better with highest fairness.

M-TCP clearly provides a comprehensive solution to the problem of hidden terminals in an ad hoc chain and cross networks, since it allows the two flows to share the bandwidth in a fair manner.

Tab. 3. M-TCP and other TCPs Fairness in [%]

	Sce. 1	Sce. 2	Sce. 3	Sce. 4
TCPW	100	50	94	86
TCP-AP	100	56	98	69
M-TCP	100	99	99	97
StandTCP	100	50	61	86

6 Conclusion

This paper studies the impact of a shared medium on TCP performance. We have proposed M-TCP, a simple algorithm to improve the fairness of flows sharing a single available bandwidth channel. Our simulations of various network topologies show that over ad hoc wireless networks, the fairness of shared flows

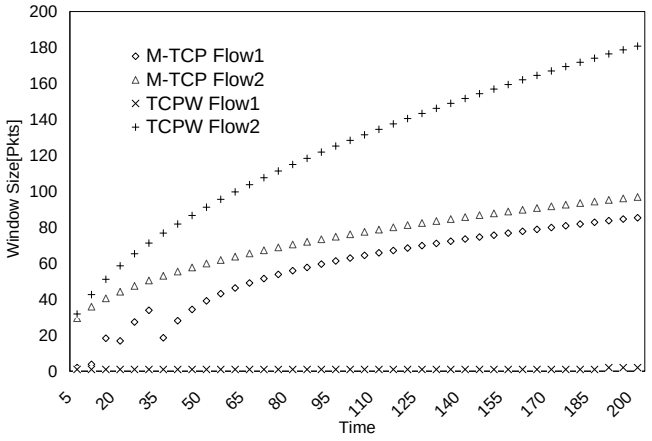

Fig. 8. M-TCP and TCPW Flows, Scenario 2

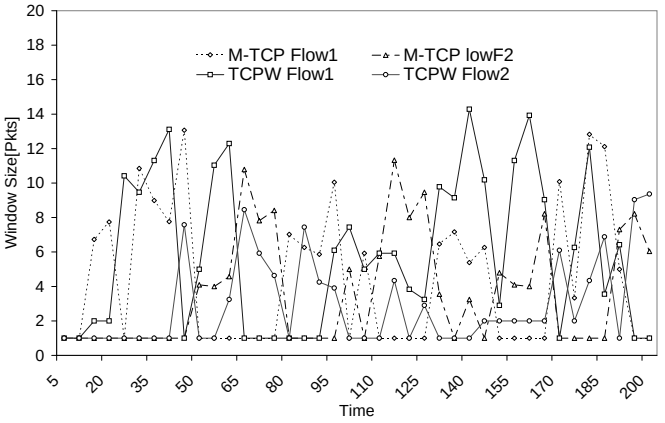

Fig. 9. M-TCP and TCPW Flows, Scenario 4

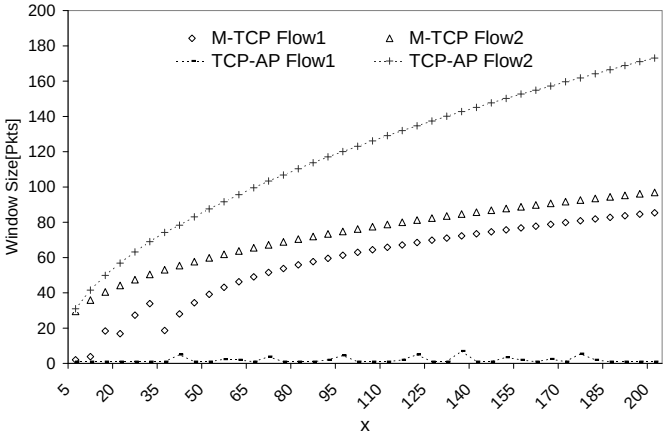

Fig. 10. M-TCP and TCP-AP Flows, Scenario 2

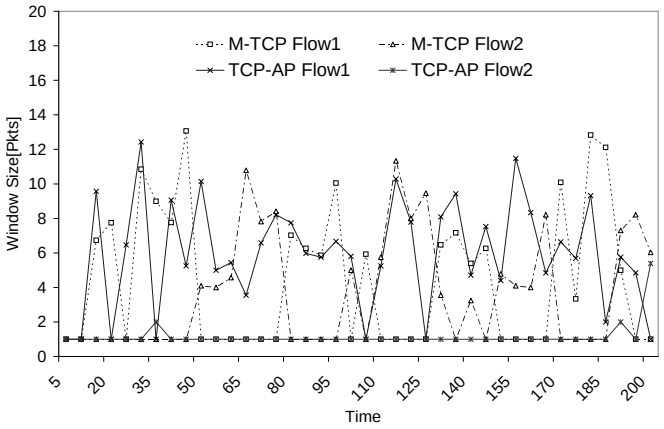

Fig. 11. M-TCP and TCP-AP Flows, Scenario 4

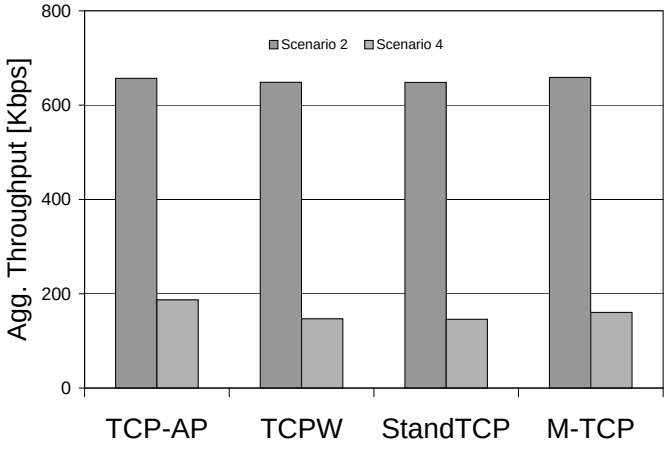

Fig. 12. Throughput of M-TCP and other TCP variants when hidden terminal exist

improves significantly if they adopt the packet transmission interval that achieves the highest possible spatial channel reuse. In all scenarios it is possible to achieve a high level of fairness while maintaining reasonable throughput and minimizing loss. We also studied the throughput as a function of the transmission interval and the TCP algorithm used. We found that M-TCP can improve throughput by 1.4% in a chain network and by 5% in cross network, with respect to standard TCP and Westwood.

References

1 **Jain R, Chiu D, Hawe W**, *A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Systems*, DEC Technical Report DEC TR-301 (1984).

2 **Fu Z, Zerfos P, Luo H, Lu S, Zhang L, Gerla M**, *The Impact of Multihop Wireless Channel on TCP Throughput and Loss*, INFOCOM'03 (2003).

3 **Fullmer CL, Garcia-Luna-Aceves JJ**, *Solutions to Hidden Terminal Problems in Wireless Networks*, In Proc. ACM SIGCOMM'97 (1997).

4 **ElRakabawy S, Klemm A, Lindemann C**, *TCP with Adaptive Pacing for Multihop Wireless Networks*, In ACM MobiHoc 2005 (2005).

5 **Liu J, Singh S**, *ATCP: TCP for mobile ad hoc networks*, IEEE J–SAC **19** (2001), no. 7, 1300–1315.

6 **Saverio Mascolo, Claudio Casetti, Mario Gerla, Sanadidi MY, Ren Wang**, *TCP westwood: Bandwidth estimation for enhanced transport over wireless links*, Mobile Computing and Networking, 2001, pp. 287-297.

7 **Li Z, Nandi S, Gupta AK**, *Achieving MAC Fairness in Wireless Ad hoc Networks using Adaptive Transmission Control*, In proc. ISCC 2004 **1** (2004), no. 28, 176–181.

8 **Perkins C, Das S**, *Ad hoc On-Demand Distance Vector Routing*, RFC 3561 (2003), available at `http://www.rfc-archive.org/getrfc.php?rfc=3561`.

9 **Xu K, Gerla M, Qi L, Shu Y**, *Enhancing TCP fairness in ad hoc wireless networks using neighborhood RED*, In Proc. ACM MobiCom (2003).

10 **Singh A, Ramanathan P**, *Spatial reuse through adaptive interference cancellation in multi-antenna wireless networks*, In Proc. GLOBECOM'05 (2005).