

Deterministic Service on IEEE 802.11-like Networks

Mario Baldi, Marco Papa Manzillo

Dipartimento di Automatica e Informatica, Politecnico di Torino

Corso Duca degli Abruzzi, 24 Torino, Italy

Email: {mario.baldi,marco.papa}@polito.it Phone: +39-011-564-{7067,7048}

Abstract

Wireless IEEE 802.11 networks are typically deployed as a convenient and flexible way to access wired backbones. However, due to the inherently shared nature and physical properties of their communication medium, performance limits are easily reached. Some standard medium access control (MAC) solutions trade simplicity for performance yielding an overall throughput possibly much lower than the actual transmission rate. Moreover, given that many applications with real-time requirements, such as VoIP phones, IPTV, videoconferencing, distributed gaming, are getting increasingly deployed and important to users, it is mandatory for the wireless access to support them. However, the delay and jitter introduced by current IEEE 802.11 standard MAC solutions in some critical, although not uncommon, scenarios can be unacceptable for real-time applications.

This work proposes a novel MAC solution for wireless networks that can offer deterministic quality of service for priority traffic while fully supporting statistical multiplexing and possibly increasing the aggregated network throughput. Being based on IEEE 802.11 physical layers and integrating some of the IEEE 802.11 MAC mechanisms, the implementation can leverage on existing components to take advantage of their economy of scale. The throughput and delay achieved with the proposed solution are compared with the ones of standard IEEE 802.11 protocols in various simulation scenarios.

1 Introduction and Motivations

There are two main issues in providing deterministic service, or quality of service (QoS), in IEEE 802.11 wireless networks: heterogeneous bit rate among network nodes and collisions. Both are not predictable and can cause possibly severe delay. The first issue stems from stations choosing the transmission bit rate based on signal quality, i.e., the signal to noise ratio (SNR). The IEEE 802.11 standard provides multiple types of modulation, each one character-

ized by different bit rate and resilience to noise. Each node chooses autonomously the best trade off between resilience and bit rate. As it will be explained later, all coordination functions specified by IEEE 802.11 medium access control (MAC) use the Distributed Coordination Function (DCF) that is based on Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). This may be critical as a transmitting node does not release the channel until a whole packet is sent. The standard includes an optional DCF fragmentation mechanism, but it is usable only to minimize the transmission error consequences and all fragments are sent consecutively. Thus, this mechanism does not prevent long transmissions from monopolizing the channel and delaying subsequent packets. This issue is exacerbated when one or more stations are using a lower transmission bit rate thus delaying transmissions by stations of the same BSS using a higher bit rate. Due to the combined effect of adaptive bit rate and deployment of CSMA/CA, as the signal to noise ratio varies in an unpredictable way for every network station (e.g. because of interference or moving away from the access point), so does the delay in accessing the BSS channel. This is exemplified by the three scenarios depicted in Figure 1 where two stations in the same BSS transmit constant bit rate streams. In first scenario (a) both nodes use a 54Mbps modulation and access to a shared channel does not impact the delay experienced by the traffic of each station. In the second scenario (b) node α has changed its modulation lowering its bit rate to 22Mbps; although node β keeps its 54Mbps modulation — hence in principle no service deterioration be expected — its packet experiences a delay D_1 resulting from node α taking more than twice as before for transmitting its packet. In the third scenario (c), where node α has lowered its transmission bit rate to 11Mbps, the time required by node α to send the first packet is so long that the transmission of the second packet is delayed. Hence although both α and β are transmitting at a bit rate higher than the bit rate of the stream each of them is sending (much higher, as far as node β is concerned), because they are using a shared channel they both find it congested and will eventually experience packet losses. The solution presented

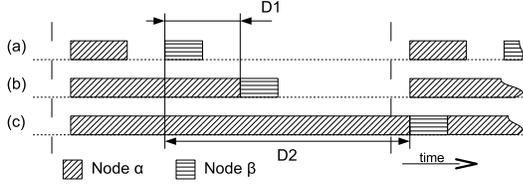


Figure 1. Variable node speed transmission

in this paper aims at avoiding such undesirable scenario where packets are lost by stations notwithstanding the good quality of their signal.

Collisions are an issue with the IEEE 802.11 MAC standards based on CSMA/CA, i.e., DCF and Enhanced Distributed Channel Access (EDCA). Since a collision is detected only at the end of a transmission (because the corresponding ACK frame is not received), when two or more stations transmit at same time (*i*) none of the transmission is successful and (*ii*) the channel is locked for the time needed for the transmission of the packet requiring most time. This will delay all the packets involved in the collision and those waiting to gain access to the channel for a time that depends on the worst case combination of packet length and selected transmission bit rate. In summary, in a BSS with many stations and a significant traffic level, the number of collisions is high, with one packet possibly colliding multiple times before being successfully transmitted, which causes unpredictable, possibly long, delays.

Both the above issues also decrease the network throughput that turns out to be significantly lower than the nominal transmission bit rate for two reasons¹: (*i*) during the time spent unsuccessfully trying to transmit colliding packets, no data is transferred through the channel and (*ii*) nodes transmitting at a lower bit rate transfer a lower amount of data in the time unit.

Both throughput reduction causes are exacerbated by the fact that according to the IEEE 802.11 DCF and EDCA standard, with the decentralized CSMA/CA the probability to gain the channel control is independent by node bit rate, but a node transmitting at a lower speed keeps the channel control longer. Hence, a small percentage of nodes transmitting at lower bit rate might have a major impact on the throughput.

The decentralized CSMA/CA algorithm is widely deployed in both infrastructure and ad-hoc BSSes. However, in infrastructure BSSes it would seem more reasonable to rely on the central role that the access point anyway has in

¹Other factors, independent of the above two issues, contribute to the throughput being lower than the line rate: frame preamble, the header being transmitted at the base transmission rate, exponential back-off time, time spacing between transmissions as specified by the IEEE 802.11 standard.

the communication² for an effective coordination of channel access. This could enable avoiding the various efficiency issues of CSMA/CA ranging from the need for inter frame spaces, to collisions – both wasting transmission time, hence lowering network throughput – to the hidden terminal problem – causing retransmissions unless additional signaling is deployed, i.e., Request To Send (RTS) and Clear to Send (CTS) messages. Along this line the IEEE 802.11 standard also specifies two coordination functions not based on CSMA/CA: Point Coordination Function (PCF) and HCF (Hybrid Coordinator Function) Controlled Channel Access (HCCA). These algorithms and DCF are not mutually exclusive: the DCF framework includes PCF or HCCA intervals. The transmission of packets previously initiated might delay the beginning of PCF or HCCA interval, thus resulting in a non deterministic behavior.

This paper proposes and assesses a synchronous channel access method called *time-driven coordination function* (TCF) that enables deterministic delay for real-time applications, while using a centralized shared channel coordinator for best-effort traffic. In essence, deterministic QoS and higher throughput are obtained by changing the IEEE 802.11 MAC layer, while keeping its widespread and inexpensive physical layer. The centralized access control avoids collisions (a cause of both delay and decrease of throughput) and ensures fair allocation of the channel from the time point of view – thus eliminating the phenomenon exacerbating the negative effects of collisions and low transmission rates on delay and throughput. Although decentralized approaches have the advantage that they can be deployed both on infrastructure and ad-hoc networks, they are not efficient in the former scenario, which is most common in nowadays access networks. Moreover, as stated in [4], also multi-hop wireless networks can be seen as a chain of BSSs, i.e., infrastructure sub-networks.

Section 2 elaborates more on IEEE 802.11 issues and presents other solutions up to now proposed in the literature. The MAC proposed by this paper is described in Section 3 while Section 4 reports some simulations results showing the benefits stemming from the proposed solution. Conclusions are drawn in Section 5.

2 Related Work

Various solutions have been proposed to support real-time applications. The IEEE 802.11e amendment specifically designed for controlled delay service, specifies two coordination functions, EDCA and HCCA. The former being based on CSMA/CA, suffers from collisions, needs large inter frame spaces, and only partially solves issues stemming from the deployment of heterogeneous transmission

²Each frame is transmitted to the access point that broadcasts it back to all the stations.

bit rates. Each packet can be assigned one of a limited set of priorities; packets of a given priority cannot be larger than a threshold defined by the BSS access point. Although this somehow improves the situation as the delay introduced by the transmission of prioritized packets is limited, it does not provide per se a deterministic service as the mechanism is not applied to best-effort packets hence a priority packet being ready for transmission after the beginning of the transmission of a non-priority packet experiences an unpredictable delay. Since the same happens also when the HCCA coordination function is deployed, it, as well, cannot guarantee deterministic services.

The Virtual MAC (VMAC) [9] technology estimates empirically the bandwidth availability simulating spurious packet deliveries but keeping listening the radio channel instead of actually send them, in such way can be evaluate the channel congestion and the probability of collision. This mechanism can improve the performances because allows to transmit data when, statistically, is less probable that a congestion can occur but can not guarantee QoS. According to the Distributed Bandwidth Allocation/Sharing/Extension protocol (DBASE) [8] priority flows are sorted in a list whose order will be followed by nodes willing to transmit, thus avoiding collision between priority flows. While a node having to send priority data waits a Point Inter Frame Space (PIFS) as defined in the IEEE 802.11 standard, before beginning its transmission, a node having to send a best-effort packet waits a longer Distributed IFS (DIFS). This algorithm allows compatibility with legacy not-priority-aware nodes, but does not solve the heterogeneous bit rate issue. Distributed Weighted Fair Queuing (DWFQ) [1] and Distributed Deficit Round Robin (DDRR) [7] are used to give different priority to delay sensitive flows. This offers a statistical improvement on the probability of having collisions, but does not offer guaranteed QoS and does not solve the heterogeneous bit rate issue.

The Multi-hop Access Collision Avoidance with Piggyback Reservations (MACA/PR) [6] is a synchronous system that reserves temporal windows for priority flows, thus avoiding collisions, but does not have a deterministic behavior when nodes change their transmission speed and is focused only on constant bit rate traffic. Distributed end-to-end Allocation of time slots for REal-time traffic (DARE) [2] is a synchronous technology for making end-to-end reservations in multi-hop scenarios thus guaranteeing a deterministic quality, but only when all stations in a BSS use the same transmission bit rate as it does not have specific support for heterogeneous bit rates. In TCF a node lowering its transmission rate might affect its own service, but other nodes keep receiving the deterministic service guaranteed by their reservation.

Another class of technologies is based on Time Division Multiple Access (TDMA). This channel access method is

based on a cyclic time structure divided in time slots. Nodes can use the shared medium only in reserved time slots and this is applied to both cyclic and bursty traffic. This can be critical for multiple reasons:

- In systems like the GSM mobile networks, time slots are reserved on a per call basis: once a reservation is done it is used for a large number of cycles. Moreover, if a node does not use one of its slot, it cannot be re-used by any other traffic. This approach is suitable for cyclic traffic while can be both inefficient or restrictive for bursty traffic.
- In systems like MAC-RSV [3], Hiperlan/2 and IEEE 802.16 reservations are done dynamically for both priority and non-priority traffic, i.e., stations can request time slots according to their transmission needs at beginning of each time cycle or at beginning of each transmission flow or when the amount of traffic they generate changes. This may be critical as a large number of reservation requests are sent in a small amount of time: non-priority requests may be served first in which case priority traffic may receives a worse service than non-priority one. When reservation requests are sent at beginning of each time cycle this misbehavior can happen for every cycle. When requests are sent at beginning of each transmission flow or when vary the load, this issue can happen when the noise level vary, so all stations have to change their modulation and renew their reservation.

These issues are not present in TCF because the TDMA approach is used only for priority traffic, while part of the rest of the traffic can be transmitted without a reservation.

3 Time-driven Coordination Function

The Time-driven Coordination Function (TCF) is here proposed to solve the aforementioned issues. Reservations granted by the access point, thus in a centralized way, eliminate issues related to collisions. In order to overcome problems stemming from heterogeneous transmission bit rates, the reservation unit is a time interval, which is independent of the amount of information to be transmitted, thus avoiding slow nodes do not keep channel control too long. Consequently, TCF operation is based on a specific time structure described in the following.

3.1 Time structure

As shown in the Figure 2, time is divided in *time cycles* that in turn are split in *Time Frames* (TFs) that are the TCF reservation unit. The first TF of the time cycle is a control time frame, called *start TF* or *s-TF*, that contains a beacon

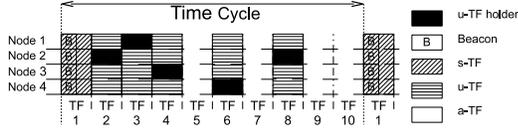


Figure 2. Time division

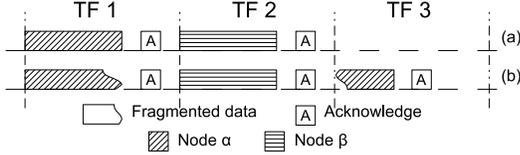


Figure 3. Variable node speed scenario

frame and control data (e.g., authentication frames). During s-TFs, nodes notify the access point of their intention to transmit a priority flow and then receive the identifiers of one or more TFs that have been reserved to them. Reserved TFs are called *u-TFs* while available ones are called *a-TFs*.

The node that holds the reservation on a u-TF can start transmitting as soon as the TF begins, i.e., without having to wait any inter frame space. According to the IEEE 802.11 standard each correctly received packet must be followed by an acknowledgement frame; hence a u-TF should enable the transmission of a packet and the corresponding acknowledgement. If the capacity of a u-TF is not sufficient the packet is fragmented. If a node has to change its modulation and thus its bit rate, and it needs more (or less) u-TFs than it currently holds, it makes a new reservation request.

In order to ensure compliance with the IEEE 802.11 hardware components, an interframe space is needed between the packet, or fragment, and the acknowledgment, and between acknowledgment and the beginning of the next time frame. The duration of inter frame spaces must be at least a Short Inter Frame Space (SIFS) as defined in IEEE 802.11 as needed by the physical layer for switching the radio interface from receiving to transmitting mode or vice versa. Finally, in order to avoid reserving whole TFs to low bandwidth flows generated by nodes transmitting at high bit rate, a TF can be divided in *sub-TFs* that can be used as reservation unit. u-TFs or portions thereof that are not used to transmit packets of the flows for which they had been reserved, as well as a-TFs, can be used for transmitting non-priority data for which a reservation is not performed, as explained in Section 3.2.

The Figure 3, shows how the issue in Figure 1 is solved with TCF. In scenario (a) both nodes use a modulation with a bit rate high enough to allow the transmission of one packet per TF: TF 1 is reserved to node α while TF 2 to

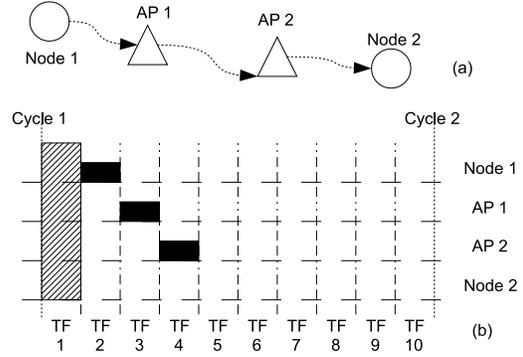


Figure 4. Multi-hop scenario, (a) topology (b) TF reservations

node β . In scenario (b), α lowers its bit rate, which in IEEE 802.11 would delay the packet from node β , as shown in Figure 1. Instead, with TCF node α realizes that its current reservation is not sufficient and makes a new request, which results in a delay while waiting for the TF 3 to be granted, but only for node α .

The time structure deployed by TCF is the same as the one required by *Time-Driven Priority* (TDP) [5], a scheduling algorithm that implements *pipeline forwarding* of packets across multiple hops with deterministic service. Hence, TCF can be straightforwardly combined with TDP to provide deterministic service in wireless multi-hop networks provided that the reservations on the various wireless links are properly scheduled to ensure pipeline forwarding. Figure 4 shows a sample scenario and reservation that enables each access point to perform *immediate forwarding* [5] of packets along the path from Node 1 to Node 2.

3.2 Unreserved Traffic

The periodic reservation solution described in Section 3.1 can be easily applied only to traffic with a predictable profile (e.g., multimedia flows), but it is not efficient with bursty traffic. Moreover, a best-effort service might be sufficient for traffic share that does not require a specific service level. Hence, packets that exceed the reservation for the flow they belong to, non-priority packets, and any packets for which a reservation is not set-up, are transmitted, as described below, as *non-priority traffic*. Transmission of non-priority traffic, however, shall not affect the service provided to packets holding a reservation to receive deterministic service, i.e., it should not increase the delay or jitter, or cause loss. This can be achieved by (i) transmitting non-priority packets during a-TFs or unused (parts of) u-TFs after the node holding a reservation on a u-frame has finished transmitting its data and (ii) ensuring that the

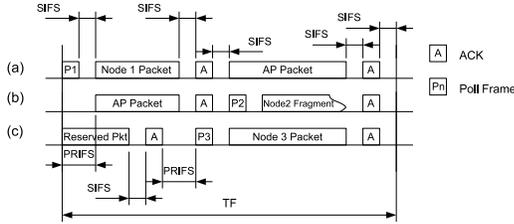


Figure 5. Non-priority traffic (a) sent during an a-TF (b) sent during an unused u-TF (c) sent in a partially used u-TF

transmission of non-priority packets does not last beyond the end of the the TFs during which it begins.

A distributed access control solution could be deployed for the transmission of non-priority packets. For example, either DCF or EDCA could be used during a-TFs if fragmentation support were introduced for avoiding transmitting across TF boundaries, while they require another minor modification for transmission during u-TFs to ensure that the node holding the reservation has a chance to transmit before any other node. Specifically, stations must wait before trying to transmit in order to make sure that the node holding the reservation is not transmitting. Afterwards CSMA/CA can be used to access the channel. For this reason, a new PRiority Inter Frame Space (PRIFS) is defined as:

$$PRIFS \geq aAirPropagationTime + aCCATime + aSIFSTime \quad (1)$$

Where $aAirPropagationTime$ is the maximum time needed for signal propagation in the air, $aCCATime$ is the time required by the physical layer to identify that the channel is busy, and $aSIFSTime$ is the duration of a SIFS. These times are related to the physical layer and are specified in corresponding standard documents (e.g. 802.11a, 802.11b). However, in case the well known hidden terminal problem arises between the node holding the reservation on a u-TF and one or more other nodes willing to transmit non-priority packets, a collision might take place, thus compromising the provision of a deterministic service. Consequently, in network scenarios where the hidden terminal problem cannot be ruled out a priori, a distributed MAC cannot be deployed during unused (portions of) u-TFs unless the RTS/CTS mechanism is deployed, which results in reduced channel utilization.

Figure 5 shows how PCF can be used to regulate channel access for non-priority traffic in the context of TCF. In the sample scenario (a), the access point sends a poll frame immediately at the beginning of an a-TF. A SIFS separates

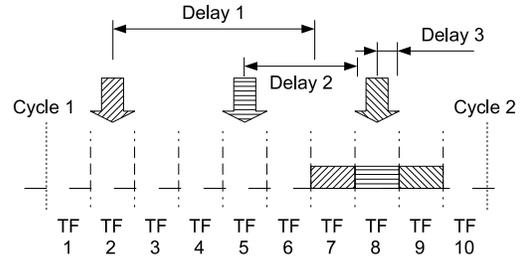


Figure 6. Delay introduced on priority packet transmission when MAC and application layers are not synchronized.

the transmission of the non-priority packet from the polling frame, which is needed by the IEEE 802.11 physical layer for switching the radio interface from receiving to transmitting mode and vice versa. During u-TFs, after having waited a PRIFS to ensure that the node holding a reservation does not have any more packets to transmit, the access point can start sending poll frames. In scenarios (b) and (c), after waiting for a PRIFS, the timing is almost the same as in scenario (a). Moreover in scenario (b) the transmitter fragments the packet to avoid the transmission to exceed the TF boundary. Figure 5 shows a possible approach where the access point alternates one packet downstream packet with an upstream packet, which is the solution deployed in the simulations presented later in the paper. .

3.3 Shortcomings and Expected Benefits

Often, priority traffic is generated with a predefined, regular pattern, (e.g., a voice over IP (VoIP) phone might transmit one packet every 40 ms). Although this fits to a MAC layer periodic reservation, if it is not synchronized with the application, a delay up to the time cycle duration is introduced before the transmission of packets. This is exemplified in Figure 6 showing a time cycle where reserved TFs are shaded and arrows represent a packet being sent by an application. Each shading pattern corresponds to a flow, i.e., a TF is reserved for packets represented by arrows with matching pattern. In a wireless multi-hop scenario in which TDP forwarding is deployed, as exemplified by Figure 4, such delay is introduced only at the first hop, while the following hops introduce a delay of only one or few TFs, depending on whether immediate or non-immediate forwarding [5] is being implemented.

The delay resulting from standard IEEE 802.11 MAC algorithms might be shorter than the one due to TCF on the first hop, especially under low load conditions. However, it is unpredictable and when the network is heavily loaded it might be even in the order of seconds, possibly at each

node in a wireless multi-hop scenario. Section 4 compares such access delay with the one achieved by TCF in various simulated network scenarios in order to assess whether the benefits provided by TCF in high traffic load conditions justify the high access delay when the application cannot be synchronized to the MAC layer (which might be impractical for existing applications).

Among the benefits expected from TCF and evaluated in Section 4 is an increased BSS throughput due to (i) collision free distributed medium access during u-TFs, (ii) avoidance of low-efficiency mechanisms such as the exchange of RTS/CTS frames during u-TFs to cope with the hidden station problem and (iii) reduction of the time interval between frames to the SIFS, during both u-TFs and a-TFs (when a centralized MAC algorithm is deployed).

Finally, it is worthwhile highlighting that while minor modifications are required at existing IEEE 802.11 MAC layers, as discussed in Section 3.2, TCF is independent from the physical layer, hence fully compatible with all existing IEEE 802.11 physical layer standards. Consequently, TCF capable interfaces could be derived by existing IEEE 802.11 interfaces with a limited effort.

4 Performance Evaluation

As initially discussed, the main motivation for the definition of TCF is providing a service with deterministic quality, specifically in BSSes where nodes use heterogeneous bit rates. This section aims at offering a first evaluation of the performance of TCF by showing that the delay experienced by priority traffic is independent of the rest of the traffic being transferred through the BSS and the bit rate of the other nodes. Moreover, the throughput achieved by the BSS is also studied.

This first TCF evaluation is based on a set of simulations performed with the widely deployed NS-2 simulator. Since TCF relies on existing functionalities and protocols of the IEEE 802.11 standard, which is well supported within NS-2, its implementation has been simple and fast, requiring roughly 3000 lines of C++ code. TCF has been extensively compared with three standard MAC technologies: IEEE 802.11 Distributed Coordination Function (DCF), the most commonly deployed nowadays, IEEE 802.11e Enhanced Distributed Channel Access (EDCA) and IEEE 802.11e HCF (Hybrid Coordinator Function) Controlled Channel Access (HCCA), all implemented in NS-2.

Simulations are run on a single BSS access network, as wireless multi-hop networks are outside the scope of this paper. In all simulations the time cycle is composed of 10 2000 μ s TFs; since the application is assumed not to be synchronized to the MAC layer, packets experience a channel access delay of up to 20 ms. Various simulation scenarios

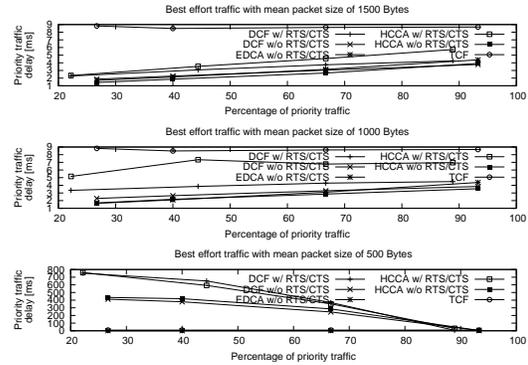


Figure 7. Access delay for priority traffic with non-priority traffic sent at 54Mbps line rate

and the corresponding results are discussed in the remainder of this section.

4.1 Throughput Assessment

The first set of simulation runs aims at estimating the maximum throughput achieved with the various MAC technologies. Ten stations greedily transmit UDP messages through the access point with a line rate of 54Mbps. Table 1 shows the aggregate throughput on the BSS for various simulation runs, each characterized by the deployment of a specific MAC solution and a IEEE 802.11 frame size of either 1500, 1000, and 500 bytes. In the scenarios where TCF is deployed, all TFs are unallocated. For the sake of completeness, the IEEE 802.11 MAC algorithms are operated both with and without the RTS/CTS mechanism enabled – except for EDCA as the RTS/CTS mechanism is not included in its NS-2 implementation. However, given that with TCF stations are not affected by hidden terminals, the comparison is most significant only with the cases where the RTS/CTS mechanism is deployed. As it can be seen in Table 1, the throughput achieved with TCF is comparable to the one achieved with other solutions when the RTS/CTS mechanism is disabled, but significantly higher (almost double in some cases) than what can be obtained when the other solutions immune from the hidden terminal problem.

4.2 Variable traffic composition scenario

This scenario is composed of an access point and ten nodes all sending non-priority traffic and some of them transmitting also priority traffic. All nodes use a 54 Mbps modulation. The total load is kept constant at 7.5 Mbps when RTS/CTS mechanism is disabled and at 4.5 Mbps when the RTS/CTS mechanism is deployed. Please note

MAC / Frame size	1500 bytes	1000 bytes	500 bytes
DCF w/o RTS/CTS	8542.73	8573.79	4809.03
EDCA w/o RTS/CTS	8044.65	8084.57	4492.8
HCCA w/o RTS/CTS	8525.98	8526.46	4791.35
DCF w/ RTS/CTS	4922.39	4974.49	2625.49
HCCA w/ RTS/CTS	4160.73	4212.84	2224.01
TCF	8138.78	8160.01	4160.03

Table 1. Achieved Throughput [Kbps]

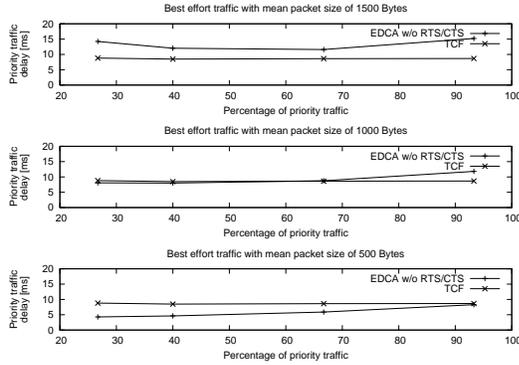


Figure 8. Access delay for priority traffic with non-priority traffic sent at a 1 Mbps line rate

that, in this scenario, TCF use the centralized approach and, thus, does not need the RTS/CTS mechanism. The percentage of priority and non-priority traffic has been varied among simulations. Priority traffic is composed of 1 Mbps flows (e.g., streaming video) and the overall amount is varied, as shown in the plots presented in this section, by changing the number of stations originating traffic flows. Non-priority traffic is composed of a fixed number of flows, one for each of the 10 wireless nodes, whose bit rate is varied to achieve the overall load specified above. Priority packet size is equal to 1000 bytes, while non-priority packet size is either 500, 1000 or 1500 bytes. Priority traffic interarrival time is constant and given by packet size divided by bit rate. Non-priority flows interarrival time is uniformly randomly distributed over an interval equal to the mean interarrival time given by the packet size divided by bit rate. In the scenarios where TCF is deployed, priority flows are transmitted during u-TFs while non-priority traffic is sent during both a-TFs and partially unused u-TFs. Figure 7 plots the average delay of priority traffic flows. Average delay with EDCA is the most stable ranging between 1.5 and 4.5 ms. When TCF is deployed delay ranges in a narrow interval between 8.5 and 8.8 ms; the variation is due to

the difference in time between the generation of a packet by an application and the beginning of the TF reserved to that application, as shown in Figure 6. This can be eliminated by introducing synchronization between MAC and application. With DCF and HCCA average delays have a very high variability (several hundreds of ms). Access delay increases as the non-priority packet size is decreased because, being the overall load constant, the number of non-priority traffic trying to access the network increases and DCF and HCCA do not succeed in keeping access delay for time sensitive traffic low.

The above issues are even more serious when nodes transmitting non-priority traffic use a slow bit rate modulation. This is shown in Figure 8 that plots the results of a set of simulations where 17 nodes communicate using a single access point, each sending a non-priority flow. Seven of the nodes can also generate a priority flow and use a 54 Mbps bit rate modulation, while the other 10 nodes transmit at 1 Mbps. The network load is the same as in the previous simulation scenario and the share of priority traffic is varied by changing the number of flows being generated. The delay experienced by priority traffic with DCF and HCCA is not shown in Figure 8 because it is very high ranging between 426.22 ms and over 10 s. Even though the delay resulting with both EDCA and TCF is acceptable delay, the former is increased when some of the nodes in the BSS use a lower bit rate, as it can be seen by comparing Figure 7 and Figure 8. EDCA issues related to the deployment of heterogeneous bit rates within one BSS are further investigated in the following section .

4.3 Heterogeneous bit rates

In order to compare the capability of the various MACs to handle variations in the bit rate used to transmit priority traffic, a set of simulations were run where three nodes send a 1 Mbps priority flow of 1000 Byte packets with an 8 ms interarrival time. One node (fast node) transmits at 54 Mbps while the transmission bit rate of the other nodes (slow nodes) is varied in each simulation run from 54 Mbps until 1 Mbps. When a node decreases its transmis-

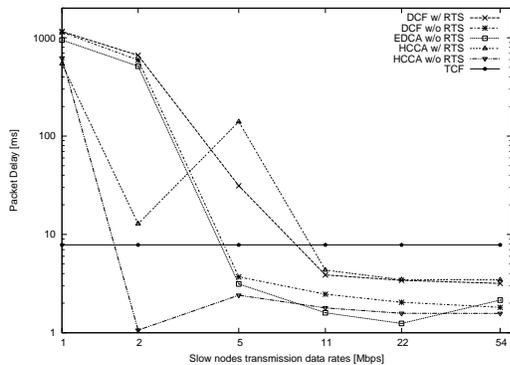


Figure 9. Fast node delay as other nodes decrease their bit rate

tion speed, it needs more time to transmit the packets of the priority flow, hence, with both HCCA and TCF the node issues a new reservation request that is handled by the access point according to the following access control policy:

- When a node N1 makes a reservation request the access point computes the necessary time T_1 and checks if enough unreserved time is available for the transmission.
- If not enough time is available, the access point searches the reservation that keeps the channel busy for largest time T_2 and the reservation holder N2 (in this scenario N2 is the node with the slowest bit rate modulation because all flows have the same bit rate).
- If $T_2 \geq T_1$ the N2 reservation is revoked and the newer request is accepted.
- Else if $T_2 < T_1$ the access point allocates all residual time to N1 but the service is not guaranteed.

This policy gives priority to faster nodes; consequently, in the presented scenario one of the two slow nodes are blocked when their transmission bit rate is 1 Mbps and 2 Mbps. As shown in Figure 9, the access delay experienced by fast node is constant if it uses TCF, but it increases considerably with all other MACs, including EDCA. In summary, TCF is the only solution that avoids that nodes with poor connectivity transmitting priority traffic, affect the service received by other nodes within the same BSS.

5 Conclusion

This paper evaluates the time-driven coordination function (TCF), a novel MAC solution for wireless networks, to support traffic requiring deterministic quality of service.

Deployed with all the IEEE 802.11 MAC solutions in various simulation scenarios, only TCF consistently offers guaranteed quality of service, in terms of access delay and jitter, even with heavy congestion and stations within the same basic service set transmitting at heterogeneous bit rates. This is achieved by deploying in a novel way synchronization among nodes, already used by IEEE 802.11 for collision avoidance algorithm, to coordinate their access to the shared medium. TCF being a MAC solution, it can be deployed together with existing IEEE 802.11 physical layers.

It is worth mentioning that, although suitable to most widespread time-sensitive applications (such as telephony, videoconferencing, and other multimedia applications), the access delay in some of the considered scenarios is higher than with other solutions. A large component of such delay is due to lack of synchronization between the application and the MAC layer, which cannot be reasonably changed in the present and near future.

However, such component is not present when TCF is deployed on a sequence of BSS in a multi-hop wireless network or wireless mesh network whose access point perform pipeline forwarding [5]. In other words, TCF is the basic building block in a promising solution, currently under study, for providing end-to-end services with guaranteed quality in wireless mesh network.

References

- [1] A. Banchs and X. Perez. Distributed weighted fair queuing in 802.11 wireless lan. *ICC*, 5, 2002.
- [2] E. Carlson, H. Karl, A. Wolisz, and C. Prehofer. Distributed allocation of time slots for real-time traffic in a wireless multi-hop network. *Proc. European Wireless*, 2004.
- [3] J. Fang and G. Kondylis. A synchronous, reservation based medium access control protocol for multihop wireless networks. *Wireless Communications and Networking, WCNC 2003. 2003 IEEE*, 2, 16-20 March 2003.
- [4] V. Gamberoza, B. Sadeghi, and E. W. Knightly. End-to-end performance and fairness in multihop wireless backhaul networks. In *MobiCom*, NY, 2004. ACM Press.
- [5] C.-S. Li, Y. Ofek, A. Segall, and K. Sohraby. Pseudo-isochronous cell forwarding. *Computer Networks and ISDN Systems*, 30(24), 1998.
- [6] C. R. Lin and M. Gerla. Real-time support in multihop wireless networks. *Wirel. Netw.*, 5(2), 1999.
- [7] W. Pattara-Aukom, S. Banerjee, and P. Krishnamurthy. Starvation prevention and quality of service in wireless lans. *Wireless Personal Multimedia Communications*, 3, 2002.
- [8] S.-T. Sheu and T.-F. Sheu. Dbase: a distributed bandwidth allocation/sharing/extension protocol for multimedia over ieee 802.11 ad hoc wireless lan. *INFOCOM*, 3, 2001.
- [9] A. Veres, A. Campbell, M. Barry, and L.-H. Sun. Supporting service differentiation in wireless packet networks using distributed control. *Selected Areas in Communications*, 2001.