Notes on WiFi scheduling and fairness

Toke Høiland-Jørgensen

toke.hoiland-jorgensen@kau.se

Introduction

These notes summarise some academic writings on WiFi scheduling and fairness issues.

The performance anomaly

When an 802.11 network contains stations running at different rates, all stations will average the bandwidth of the *slowest* station. This has been known for a while as the 802.11 *performance anomaly* [6] and stems from the fact that the 802.11 MAC gives all stations the same probability of accessing the medium, regardless of the length of their transmission. This results in *throughput fairness* between the stations. A more efficient resource allocation is *air time fairness* or *temporal fairness*, which is equivalent to *proportional fairness* between the stations [7].

In the presence of both 802.11g and 802.11n stations in the same network, the performance anomaly also occurs; but it can be somewhat mitigated by the fact that the slow stations (running on 802.11g) do not aggregate packets, while 802.11n stations can turn on aggregation to be able to transmit more packets for each transmit opportunity [13].

A way to mitigate the performance anomaly is by adjusting the transmission size according to the bandwidth. One scheme to do this uses the policy routing mechanism of Linux to cap the MTU size in an attempt to achieve a packet size ratio that matches the transmission rate ratio between stations [3]. While this scheme is practically implementable (the authors' prototype uses perl scripts on the access point side for an 802.11b network on Linux 2.4), it is limited by the small span of the possible MTU values. In addition, smaller MTU sizes increases the transmission overhead.

Jiang et al outlines two basic approaches to achieve air time fairness in pre-802.11n networks (i.e. with no packet aggregation) [7]: (1) Using the TXOP feature of the 802.11e QoS standard to allocate equal-length transmit opportunities to each station (where multiple packets can be sent in each TXOP). (2) Adjusting the size of the minimum contention window (*CWMin*) so that stations with lower transmission rates have a larger window, and thus a lower probability of accessing the medium. The authors also propose an algorithm for stations in ad-hoc networks to approximate their fair share of air time, by listening to stations within their transmission range; and then using that information to adjust its transmission parameters to achieve air time fairness. The authors validate their scheme in simulation (based on the *CWMin* adjustment approach) and find that it achieves close to perfect air time fairness.

An analysis of different ways of achieving air time fairness is given by Lin et al [10]. They classify the approaches into two categories: the deterministic approach, which simply sets the TXOP length based on the transmission rate, and the statistical approach, which modifies the *CWMin* or the inter-frame spacing to achieve different medium access probabilities for different stations. The authors note that both approaches can lead to increased latency: Changing the length of the TXOP means that when a station transmits several packets back-to-back, the other stations has to wait for longer to access the medium (inducing latency). And changing the *CWMin* and inter-frame spacing means stations will wait longer before transmitting, which also introduces delay. The authors suggest an approach that modifies *both* the inter-frame delay and the *CWMin* so as to achieve the fairness target while minimising the added delay. They then outline a machine learning approach that solves for the right parameters to achieve this.

Razafindralambo et al propose solving the performance anomaly by aggregating packets (this is also pre-802.11n) and propose a simple scheme to do so [11]. Their scheme consists of each station continuously sensing the channel and keeping track of the longest continuous transmission interval observed from other stations. Then, when it becomes a station's turn to send, it will simply keep emitting packets as soon as the previous packet is ACKed, until it has occupied the channel for the same amount of time as it has sensed other stations do. They simulate their solution and find it provides good fairness; however, it is not clear that it is compatible with non-modified 802.11 stations.

Another scheme to achieve air time fairness through scaling the contention window is presented by Joshi et al [8]. They provide an analytical model for computing the achieved throughput from 802.11b protocol parameters and uses this together with feedback on achieved error and transmission rates from the WiFi device to define a protocol that dynamically scales the contention window to achieve the desired fairness properties. They evaluate their analytical model through simulations and conclude that it achieves good fairness and reacts quickly to changes in the channel conditions.

Another way of adopting the length of each transmission opportunity to achieve fairness is by using the aggregation features of 802.11n and later. Kim et al propose a scheme to do just this [9]. The scheme describes the optimisation problem of selecting the right amount of packets (assumed to all be of the same size) to achieve a target air time for each transmission opportunity. The 802.11n standard offers two levels of aggregation, A-MSDU (which aggregates several subframes into a MAC frame with a common header) and A-MPDU (which aggregates several MAC frames into the same physical frame), which can be combined in a single transmission. There is a tradeoff in efficiency and reliability between the two aggregation methods, because each MAC frame can be ACK'ed separately, but incurs an additional header overhead. The aggregation scheme proposed in the paper uses channel loss information to select the right number of aggregation units at each level for every transmitted frame. The objective is to get as close to the reference target transmission time (for which they determine 3 ms to be a good value) as possible, subject to the transmission rate and limits on max aggregation size imposed by the standard. The paper formulates the optimal integer programming problem, as well as an efficient estimation procedure that can be solved in real time. They evaluate their scheme in simulation and find that it achieves very good fairness and higher aggregate throughput compared to the baseline 802.11 scheme.

Other scheduling approaches

Another approach to mitigating the performance anomaly that looks only at the access point scheduling is proposed by Garroppo et al [4]. They observe that in many cases, when a WLAN runs in infrastructure mode, most of the traffic originates at the access point, due to clients downloading more than they upload. This leads them to designing a scheduler for the access point that serves stations in a manner that maximises fairness. They do this by introducing a scheduling module (called the deficit transmission time (DTT) scheduler) in a Linux WiFi device driver that schedules traffic towards different stations, based on feedback from the actual transmission time of packets to each station. Their scheme features one queue per associated station. Each queue has a token bucket attached that is used to track the actual transmission time (including retransmits) of each packet towards a destination. All stations with packets queued get their tokens replenished at their fair share rate (measured in time), and when a packet is transmitted, the time taken is subtracted from that station's bucket. The scheduler consists of simply selecting the queue with the largest token balance at dequeue time (or slightly before; the scheme keeps one extra packet queued in the hardware to prevent throughput from suffering from having to wait on the driver when a transmission opportunity becomes available).

Using TDMA scheduling is another way to provide fairness. Ben Salem and Hubaux propose a scheme for TDMA scheduling in a mesh network that ensures fair bandwidth allocation to all connecting clients, no matter where in the mesh they connect [1]. Torfs and Blondia outline some of the difficulties of implementing TDMA on commodity hardware [14], evaluating the timer reliability in Linux and implementing a TDMA transmission scheme using the *ath9k* driver in access point mode.

A modification to the CSMA/CA scheme of 802.11 is proposed by Sanabria-Russo et al [12]. They outline previous suggested changes to the scheme and develop them further in their *CSMA/ECA with hysteresis and fair share* (the "E" is for "Enhanced"). This scheme uses a deterministic back-off when a packet is successfully sent, rather than the random back-off in baseline 802.11. In addition, they propose keeping the same back-off stage while the station has packets to send (this is the hysteresis), and to send a number of packets proportional to the back-off stage, thus compensating for the lower medium access probability. They evaluate this scheme in simulation and note that it can lead to completely collision-free operation in many cases, increasing the aggregate goodput. However, they do not evaluate the latency and air time fairness consequences of their scheme.

An overview of scheduling algorithms for TDMA wireless links is given by Cao and Li [2]. They outline several goals for scheduling of wireless links, and survey several scheduling algorithms in different categories. However, they note that since all the scheduling algorithms model the wireless channel by a two-state Markov chain, none of them are suitable for an CDMA-style MAC where the link can have more than two states.

Hajlaoui et al examine 802.11n aggregation and (unsurprisingly) conclude that naively aggregating packets (i.e. waiting for a full aggregate before sending) results in higher latency the larger the aggregate [5]. They propose an aggregation scheduler that uses the diffserv priorities to make aggregation decisions, deferring background and best effort traffic until a full (fixed size) aggregate can be constructed, but sending voice and video traffic immediately.

References

[1] Naouel Ben Salem and Jean-Pierre Hubaux. 'A fair scheduling for wireless mesh networks'. In: *WiMesh*. 2005. URL: `http://infoscience.epfl.ch/record/63669` (visited on 09/17/2015).

[2] Yaxin Cao and Victor OK Li. 'Scheduling algorithms in broadband wireless networks'. In: *Proceedings of the IEEE* 89.1 (2001), pp. 76–87. URL: `http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=904507` (visited on 09/17/2015).

[3] Joseph Dunn et al. 'A practical cross-layer mechanism for fairness in 802.11 networks'. In: *Broadband Networks, 2004. BroadNets 2004. Proceedings. First International Conference on*. IEEE, 2004, pp. 355–364. URL: `http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1363824` (visited on 09/16/2015).

[4] Rosario G. Garroppo et al. 'Providing air-time usage fairness in IEEE 802.11 networks with the deficit transmission time (DTT) scheduler'. en. In: *Wireless Networks* 13.4 (Aug. 2007), pp. 481–495. ISSN: 1022-0038, 1572-8196. DOI: `10.1007/s11276-006-9201-7`. URL:

http://link.springer.com/10.1007/s11276-006-9201-7 (visited on 09/11/2015).

[5] Nasreddine Hajlaoui and Ihssen Jabri. 'Experimental Performance Evaluation and Frame Aggregation Enhancements in IEEE 802.11n WLANs'. In: *Wireless and Mobile Networking Conference (WMNC), 2012 5th Joint IFIP*. IEEE, 2012, pp. 64–69. URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6416145 (visited on 09/16/2015).

[6] Martin Heusse et al. 'Performance anomaly of 802.11 b'. In: *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*. Vol. 2. IEEE, 2003, pp. 836–843. URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1208921 (visited on 09/14/2015).

[7] Li Bin Jiang and Soung Chang Liew. 'Proportional fairness in wireless LANs and ad hoc networks'. In: *Wireless Communications and Networking Conference, 2005 IEEE*. Vol. 3. IEEE, 2005, pp. 1551–1556. URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1424745 (visited on 09/11/2015).

[8] T. Joshi et al. 'Airtime Fairness for IEEE 802.11 Multirate Networks'. In: *IEEE Transactions on Mobile Computing* 7.4 (Apr. 2008), pp. 513–527. ISSN: 1536-1233. DOI: 10.1109/TMC.2007.70740. URL: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4358990 (visited on 09/11/2015).

[9] Minho Kim, Eun-Chan Park, and Chong-Ho Choi. 'Adaptive Two-Level Frame Aggregation for Fairness and Efficiency in IEEE 802.11n Wireless LANs'. en. In: *Mobile Information Systems* 2015 (2015), pp. 1–14. ISSN: 1574-017X, 1875-905X. DOI: 10.1155/2015/548109. URL: http://www.hindawi.com/journals/misy/2015/548109/ (visited on 09/16/2015).

[10] Pochiang Lin, Wu-I. Chou, and Tsungnan Lin. 'Achieving airtime fairness of delay-sensitive applications in multirate IEEE 802.11 wireless LANs'. In: *Communications Magazine, IEEE* 49.9 (2011), pp. 169–175. URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6011749 (visited on 09/11/2015).

[11] Tahiry Razafindralambo et al. 'Dynamic packet aggregation to solve performance anomaly in 802.11 wireless networks'. In: *Proceedings of the 9th ACM international symposium on Modeling analysis and simulation of wireless and mobile systems*. ACM, 2006, pp. 247–254. URL: http://dl.acm.org/citation.cfm?id=1164761 (visited on 09/14/2015).

[12] Luis Sanabria-Russo et al. 'Future evolution of CSMA protocols for the IEEE 802.11 standard'. In: *Communications Workshops (ICC), 2013 IEEE International Conference on*. IEEE, 2013, pp. 1274–1279. URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6649433 (visited on 09/11/2015).

[13] Vivek Shrivastava et al. '802.11 n under the microscope'. In: *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*. ACM, 2008, pp. 105–110. URL: http://dl.acm.org/citation.cfm?id=1452533 (visited on 09/16/2015).

[14] Wim Torfs and Chris Blondia. 'TDMA on commercial of-the-shelf hardware: Fact and fiction revealed'. In: *AEU-International Journal of Electronics and Communications* 69.5 (2015), pp. 800–813. URL: http://www.sciencedirect.com/science/article/pii/S1434841115000254 (visited on 09/16/2015).