

# Study of the Intel WiFi Rate Adaptation Algorithm

Rémy Grünblatt, Isabelle Guérin-Lassous, Olivier Simonin

► **To cite this version:**

Rémy Grünblatt, Isabelle Guérin-Lassous, Olivier Simonin. Study of the Intel WiFi Rate Adaptation Algorithm. CoRes 2019 - Rencontres Francophones sur la Conception de Protocoles, l'Évaluation de Performance et l'Expérimentation des Réseaux de Communication, Jun 2019, Saint-Laurent-de-la-Cabrerisse, France. pp.1-4. hal-02126333

**HAL Id: hal-02126333**

**<https://hal.inria.fr/hal-02126333>**

Submitted on 11 May 2019

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Study of the Intel WiFi Rate Adaptation Algorithm

Rémy Grünblatt<sup>1,2</sup>, Isabelle Guérin-Lassous<sup>1</sup>, Olivier Simonin<sup>2</sup>

<sup>1</sup>Univ Lyon, UCBL, EnsL, CNRS, Inria, LIP, F-69342, LYON Cedex 07, France.

<sup>2</sup>Univ Lyon, INSA Lyon, Inria, CITI, 6 avenue des Arts, 69621 Villeurbanne Cedex, France

---

Dans cet article, nous décrivons et analysons l'algorithme d'adaptation du débit de transmission implanté dans les cartes WiFi Intel récentes, et utilisé notamment dans les cartes réseaux des drones Intel Aero Ready-To-Fly. Les performances de cet algorithme sont ensuite évaluées expérimentalement dans un scénario simple.

**Mots-clefs :** IEEE 802.11, WiFi, Rate Adaptation Algorithm, Intel

---

## 1 Introduction

Unmanned aerial vehicles (UAVs) can help to conduct a variety of applications such as mapping, monitoring, or delivery, both in a civil or military context. Equipped with off the shelf WiFi Network Interface Controllers (WNICs), UAVs can communicate together as well as with a variety of compatible devices, such as smartphones or personal computers. While WNICs support multiple transmission rates, the IEEE 802.11 standard makes no attempt at defining algorithms to select the transmission rate. However, rate adaptation plays an important part in 802.11 performances, such as throughput or delay [AGH13, XHF12]. These performances have an impact on the communication quality but also on all the UAV operations that rely on these communications. It is therefore important to well understand the core of the used rate adaptation algorithms.

In this paper, we study the rate adaptation algorithm (RAA) of an Intel Aero Ready-To-Fly UAV, equipped with an Intel® Dual Band Wireless-AC 8260 WNIC. The contributions of the paper are:

- The extraction of the rate adaptation algorithm from the Intel WNIC driver code and its presentation for the first time;
- The experimental evaluation of this algorithm in a simple scenario.

## 2 The IWL-MVM-RS Rate Adaptation Algorithm

RAA usually acts on successful or failed transmissions to scale up or down the transmission rate, but they also often include look-around decisions used to evaluate potential gains in performance. For full-MAC drivers, the rate adaptation is done in the driver or in the firmware, while for soft-MAC drivers, it can be done in the mac80211 component of the Linux kernel, allowing a single implementation to drive different piece of hardware. Some of soft-MAC drivers still use their own RAAs, which is the case of the IWLWIFI driver used by Intel wireless chips. It comes with its own algorithms: IWL-AGN-RS and IWL-MVM-RS, the former being unmaintained, and limited to 802.11n hardware, while the latter is being used with 802.11ac compatible MVM hardware, and is currently undergoing changes to support 802.11ax.

**Algorithm overview** IWL-MVM-RS takes care of managing the Modulation and Coding Scheme (MCS) index, but also whether to transmit in a legacy mode (802.11a or 802.11g) or in a non-legacy mode (802.11n or 802.11ac) in a SISO or MIMO way. It chooses which antenna or subset of antennas to transmit with, and whether a Short Guard Interval (SGI) or a Long Guard Interval (LGI) is used. It decides when to enable frame aggregation.

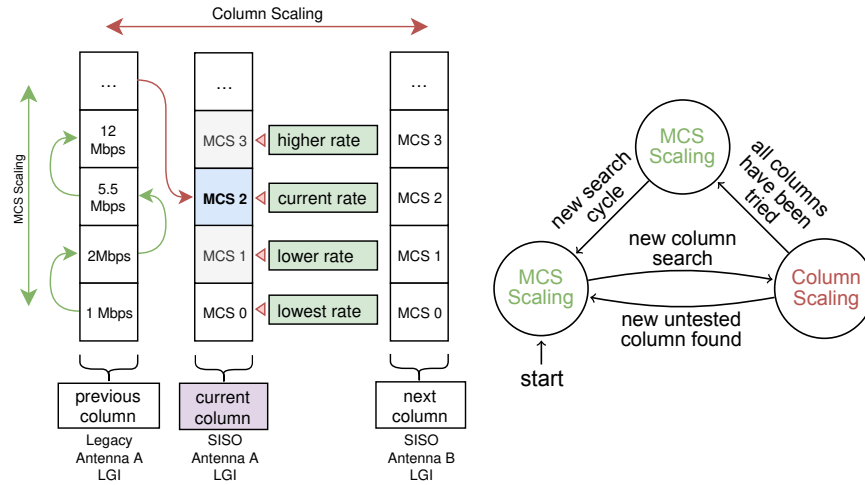


Fig. 1: **(Left)** Example of sequence of decisions made by the RAA; **(Right)** Flowchart of the different states of the RAA IWL-MVM-RS.

IWL-MVM-RS has two main components: MCS Scaling<sup>†</sup> and Column Scaling. MCS Scaling tries to maximize the throughput by only changing the MCS, while Column Scaling tries to find a better *column*, which is a combination of *mode* (legacy, SISO, MIMO), *guard interval*, and *antenna configuration* parameters. The algorithm starts with the lowest parameters (associated with the worst throughput, but with the best reliability) and interleaves MCS Scaling phases and Column Scaling phases, forming a "search cycle". Column Scaling starts when the MCS Scaling phase chooses not to change MCS. The alternation of MCS Scaling and Column Scaling continues until all the columns have been tried, which means the end of the search cycle and the beginning of a new one. Figure 1 sums up the different steps of the algorithm, steps described in more details hereafter.

**MCS Scaling** The MCS scaling algorithm can take one of the following decisions: lowering the MCS, raising the MCS, or keep using the current MCS. The decision is made in a deterministic manner, according to the maximum theoretical throughput and the measured throughput of adjacent MCS indexes as well as the measured throughput of the current MCS. It therefore prevents from switching to a MCS implying a theoretical throughput lower than the current measured throughput, or to a MCS whose measured throughput is worse than the current one, for example.

The theoretical throughput is hardcoded into tables for each mode (legacy, SISO, MIMO), each MCS index, and for the four possible guard interval and aggregation parameters (SGI, LGI, SGI+AGG, LGI+AGG). The measured throughput for each MCS index is computed by multiplying the success ratio of up to the last 62 frame transmissions at this MCS (with at least 8 successful transmissions or 3 failed transmissions) with the theoretical throughput when using this MCS. The decisions are the following:

1. if the success ratio is too small ( $< 15\%$ ) or the measured throughput is zero, decrease the MCS index;
2. else, if the measured throughput with the higher adjacent MCS index is better than the measured throughput of the current MCS, or unknown, increase the MCS index;
3. else, if the measured throughput with both the lower adjacent and higher adjacent MCS indexes are worse, or if success ratio is big enough ( $> 85\%$ ), maintain the MCS;
4. else, decrease the MCS index.

**Column Scaling** Each column has a set of "next columns" that the driver will try if they can theoretically beat the current measured throughput (by looking at the theoretical throughput of the columns), and if they have not been already tested. When trying a new column, if the measured throughput in this column is better than the throughput in the previous one, the RAA keeps using it. Otherwise the column is marked to avoid trying it again during the search cycle, and the RAA reverts back to the old column. The initial starting MCS

<sup>†</sup> Even if the MCS concept does not exist until 802.11n, this term is used as a handy shortcut to refer to both MCS and data rates.




Fig. 2: Test Setup. The dashed parts are the parameters: AP movement and distance between AP and STA.

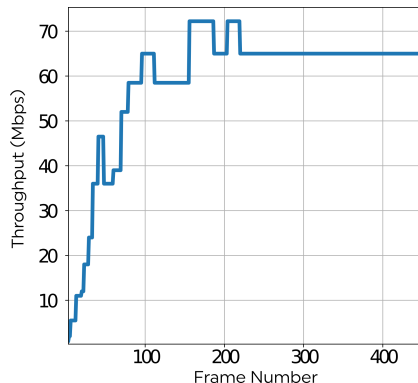


Fig. 3: Example of transmissions rates chosen when starting. This pattern is typical of IWL-MVM-RS.

index in the new column is chosen according to the success ratio: if it is high enough (more than 85%), the smaller MCS index whose theoretical throughput is higher than the current theoretical throughput is chosen. Otherwise, the smaller MCS index whose theoretical throughput is higher than the current measured throughput is chosen. After a column switch, measured throughput of the previous column are dropped.

**New Search Cycle** After the end of a search cycle, the algorithm does MCS Scaling until the start of a new search cycle. This start is triggered when:

1. too many frames failed (160 in legacy, 400 otherwise) since the beginning of the previous cycle;
2. too many frames succeeded (400 in legacy, 4500 otherwise) since the beginning of the previous cycle;
3. too much time has been spent after the end of the previous search cycle (5 seconds);

### 3 Experiments

UAV establishing a communication network can position themselves according to a variety of parameters (e.g. distance, RSSI). To determine which parameters are relevant in order to maximize the end-to-end (E2E) throughput, we conduct an experiment whose setup is shown in Figure 2.

#### 3.1 Experiment Description

A station (STA, playing the role of a drone) using a IWL-MVM-RS WNIC communicates using a 802.11n Wireless LAN with an access point (AP, playing the role of a controller). Both the STA and the AP are placed on a support at a height of approximately 1m, and experiments are performed outdoor in an open environment, with line-of-sight conditions. The AP is a TP-Link TL-WR802Nv4 (mt7603 system on chip) running OpenWrt (Linux 4.9.73). The STA is a Dell Precision 5520 running Archlinux (Linux 4.20.7), and is equipped with an Intel wireless-ac 8260 WNIC. A CSL AC1200 USB WNIC (Realtek RTL8812AU) is plugged into the STA and is used for monitoring purpose. The wireless cards are configured to use the channel 11 (2462Mhz, 20MHz width). Each test lasts 30 seconds and is repeated 5 times.

We study the network throughput as a performance metric: we use iPERF3 to generate a saturating UDP traffic on the STA and to measure applicative end-to-end (E2E) throughput. During a test, the distance between the AP and the STA is fixed. The AP can be put in movement by making it move around itself, without changing too much the distance to the STA. Two kinds of scenarios are studied: one where both the STA and AP are static during each test, and the other where the AP is put in movement<sup>‡</sup>. Data recorded, by using the monitoring WNIC, are then compared to the E2E throughput as reported by iPerf3.

#### 3.2 Results

Figure 3 shows the evolution of the transmission rates used on the first transmitted 450 frames. This evolution is typical of the IWL-MVM-RS algorithm. The obtained results on the the full experiments<sup>§</sup> show that

<sup>‡</sup> In this case, the tests have not been completed for all distances. However, we have decided to present the results.

<sup>§</sup> not provided here due to space limitation

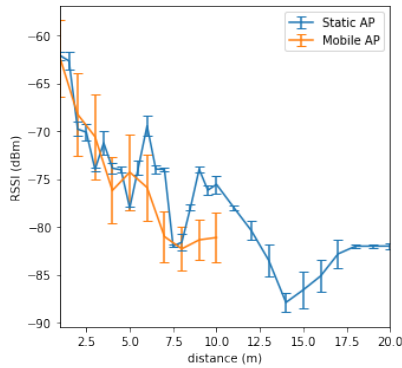


Fig. 4: Average RSSI of the AP with respect to the distance.

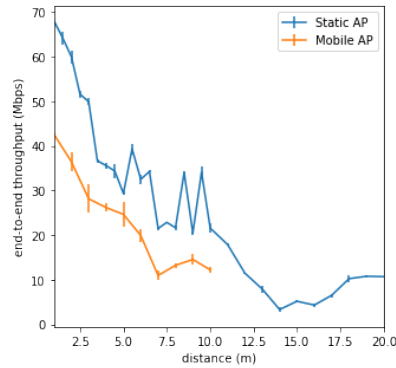


Fig. 5: Average E2E throughput with respect to the distance.

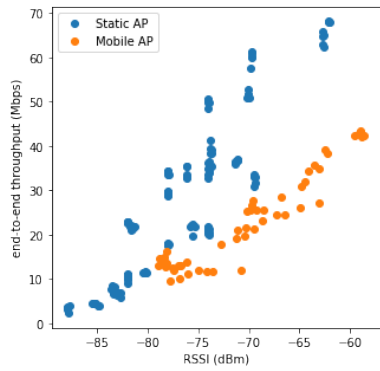


Fig. 6: Average E2E throughput with respect to RSSI.




Fig. 7: Average E2E throughput with respect to mean transmission rate.

transmission rate regularly changes and never remains constant even if the STA and the AP are fixed and whatever the distance between them. In Figures 4 and 5, we globally observe negative correlation between the RSSI and E2E throughput with respect to the distance. Mobility of the AP increases the number of transmission rate changes, leading to a reduction of E2E throughput. However, the presence of peaks and wells argues for a smart placement of UAVs in order to maximize E2E throughput since, in some cases, higher RSSI and E2E throughput are obtained on distances larger than others. Moreover, E2E throughput is loosely correlated with average RSSI, as shown in Figure 6: throughput can go from single to double for a same RSSI. On the other hand, a strong correlation is observed between E2E throughput and mean transmission rate in Figure 7. Mean transmission rate is computed as the average on the transmission rates used on all the frames sent during the experiment. Mean transmission rate therefore appears as an easy-to-compute metric which gives a good indicator of the E2E throughput in this scenario (one hop, saturated conditions).

## Acknowledgements

We would like to thanks the Direction Générale de l'Armement (DGA) and the Fédération Informatique de Lyon (FIL) for their financial support.

## References

- [AGH13] M. Asadpour, D. Giustiniano, and K. A. Hummel. From ground to aerial communication: Dissecting wlan 802.11n for the drones. In *WiNTECH*, pages 25–32. ACM, 2013.
- [XHF12] D. Xia, J. Hart, and Q. Fu. On the performance of rate control algorithm Minstrel. In *PIMRC*, pages 406–412. IEEE, 2012.