



# Architect a Next-Gen 802.11ac Wave 3 Software-Defined Modem

---

## White Paper

**Lisa Meilhac, PhD.**  
WLAN System Architect, CEVA

**Franz Dugand.**  
Sales and Marketing Director, CEVA

## Table of Contents

1.	Introduction	3
2.	Market Vision	3
3.	Product Overview	4
4.	Hardware Accelerators	6
5.	Example 1: 802.11ac 4x4, 4 stream, 160MHz configuration	10
6.	Example 2: 802.11ac 8x8, 8 Streams, 160MHz Configuration	13
7.	Conclusion	14

## 1. Introduction

CEVA's RivieraWaves connectivity platforms address the continuous need for more highly integrated System-on-Chips (SoCs) to reduce cost, power consumption and size. This family incorporates Wi-Fi and Bluetooth Smart and Smart Ready solutions, which can be standalone or combined together for total combination solutions.

The RivieraWaves Wi-Fi IP family offers a comprehensive suite of platforms for embedding Wi-Fi 802.11a/b/g/n/ac into SoCs/ASSPs. Optimized implementations are available targeting a broad range of connected devices, including smartphones, wearables, consumer electronics, smart home, gateway, industrial and automotive applications. Each of the RivieraWaves Wi-Fi platforms incorporates the Upper MAC (UMAC), Lower MAHarC (LMAC) and PHY modem functions.

In particular, the RivieraWaves Wi-Fi IP family includes a high performance Wi-Fi platform, called RivieraWaves Stream, which addresses the most advanced use cases applicable to access points, media gateways and Wi-Fi offload in small cells. Leveraging the processing power of the CEVA-XC4210 DSP for the PHY modem function, it is scalable to address up to several hundreds of users and offers flexible co-existence with LTE/LTE-A in infrastructure applications. The RivieraWaves Stream is available today for configurations up to 802.11ac 4x4.

This paper presents an evolution of the RivieraWaves Stream architecture for next-generation 802.11 ac Wave 3, which can support complex configurations up to 8x8 MU-MIMO with 160MHz bandwidth.

## 2. Market Vision

ABI Research forecasts that consumer and enterprise WLAN access point shipments will reach 171.5 million and 25 million respectively at the end of 2019<sup>1</sup>. WLAN access point shipments are today predominantly based on 802.11ac 4x4. According to another study by Machina Research, the average US home is expected to have 20 connected devices by 2020<sup>2</sup>, all of which need to interface to the same consumer WLAN access point. Enterprise WLAN access points are very likely to be connected to a much higher number of devices. The total installed base of Wi-Fi-enabled devices is expected to surpass five billion devices by the end of 2015. Single-antenna and double-antenna devices dominate that population, with smartphones making up over 50 percent of annual Wi-Fi shipments.

The above trend has driven the need for deployment of devices based on the so-called 802.11ac Wave 3 specification, which brings the following key additional benefits:

- Twice the bandwidth for twice the throughput - up to 160MHz, while 802.11ac Wave 1 is limited to 80MHz. Despite being part of the 802.11ac Wave 2 specification, many 802.11ac Wave 2 solutions still do not support 160MHz bandwidth.
- Twice the number of antennas for even higher throughput: up to 8x8 while Wave 1 and Wave 2 are limited to 4x4. 8x8 is particularly recommended for WLAN enterprise access points due to the high number of connected devices.
- MU-MIMO (Multi User MIMO). This is the most important feature introduced in 802.11ac Wave 2. In a Wave 1 solution, a 4x4 access point can talk to 1x1 client devices one after the other. This limits the throughput and network utilization to 1x1 only, hence not taking benefit of the 4x4 capability. In a Wave 2 system, a 4x4 access point can talk to up to four 1x1 clients at the same time with each of the four clients receiving the signal from one of the four spatial streams. Taking advantage of the 4x4 capability of the access point significantly improves network utilization,

RivieraWaves Stream 802.11ac targets three market segments for 802.11ac access points:

- Enterprise access points requiring up to eight spatial streams (8x8) to support high number of clients, from 1x1 up to 4x4
- Home / consumer access points - the biggest market, requiring up to four spatial streams
- Low-cost mobile access points with up to two spatial streams.

It is important to provide a scalable solution that can meet the requirements of these three markets and this paper presents a scalable solution for an 802.11ac Wave 3 modem architecture that can scale from 2x2 up to 8x8.

### 3. Product Overview

The CEVA RivieraWaves Stream 802.11ac Wave 3 software-defined modem (SDM) subsystem is an extremely high-performance Wi-Fi modem, supporting a large range of configurations up to very large MIMO dimensions.

#### 3.1. Features Set

The CEVA RivieraWaves Stream 802.11ac Wave 3 SDM modem includes all the signal processing for transmit and receive between the MAC interface and the ADC/DAC, including radio control and AGC control. It supports the following features:

- The majority of the 11ac optional modes
- All modulation schemes up to 256-QAM: MCS0-MCS9, from one to eight spatial streams

- 1024-QAM modulation (MCS10 and MCS11)
- Long Guard Interval (800ns) and Short Guard Interval (400ns)
- Space Time Block coding (STBC) for improved link reliability, minimizing the effects of scattering, reflection and refraction
- Low Density Parity Check (LDPC), which improves receive sensitivity by 2-3 dB compared to a Viterbi decoder
- Transmit beamforming, as a beamformer and as a beamformee
- Multi User MIMO (MU-MIMO)

### 3.2. Modem Subsystem Overview

The CEVA Wave 3 SDM modem subsystem is a hybrid design mixing hardwired units with CEVA-XC cores and offering the best trade-off between size and flexibility. It is provided with a reference control and processing software resulting in a complete and fully functional Wi-Fi modem.

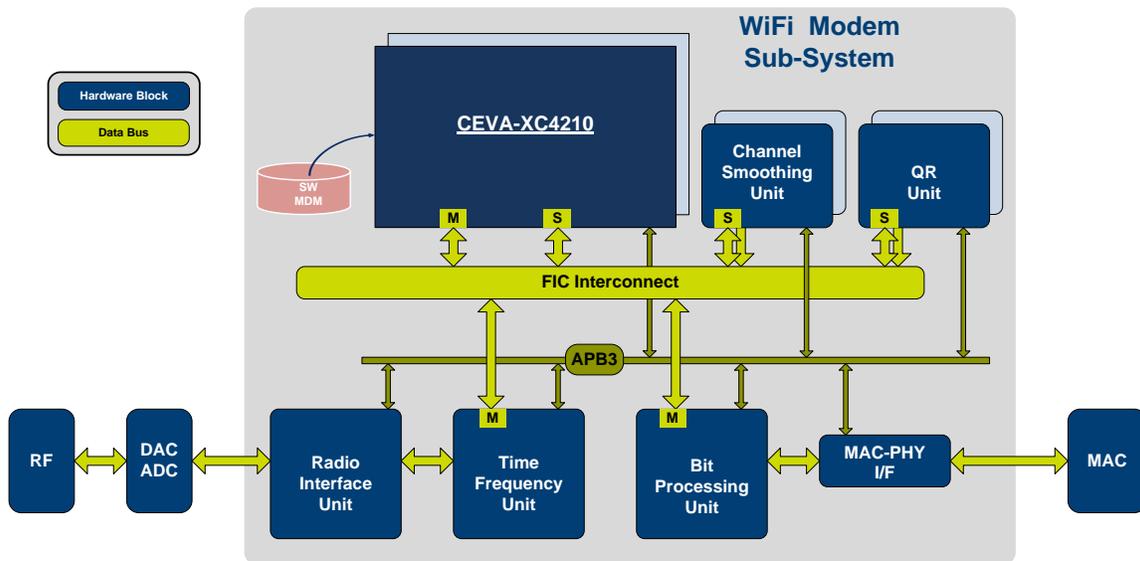


Figure 1 - CEVA Wave 3 SDM Modem Architecture

This modem should be combined with the CEVA Wave 3 Wi-Fi MAC subsystem and an analog/RF subsystem to form a complete Wi-Fi system.

Some of the modem signal processing algorithms have been selected to be implemented with dedicated logic because they are computationally demanding and are not expected to change. But all the key Wi-Fi modem algorithms are implemented on the DSP cores and benefit from its flexibility. The CEVA Wave 3 Wi-Fi modem reference software is composed of sophisticated algorithms taking full advantage of the highly

powerful vector capabilities of the CEVA-XC core. Critical algorithms are coded in pseudo-float, minimizing the implementation loss and ensuring the best performance.

The software communicates with the hardwired units using registers and interrupts. Data exchanges between hardwired units and the core are accomplished via a dedicated memory-based interconnection using the DSP core's data memory.

This solution represents a highly scalable platform, with the number of cores as well as the hardwired unit dimensions depending on the targeted configuration. The CEVA-XC can run at high frequencies and most configurations can be handled by a unique core while still leaving a large percentage of the core processing power free. Customers can make use of this processing headroom for differentiation, extensions or enhancements.

For the largest configuration, such as 8x8, the architecture is scaled by using two CEVA-XC cores to share the processing load of the critical elements. Corresponding to the frequency domain processing of the preamble and data fields, the split is very simply achieved: each core processing half of the subcarriers. The Time Frequency Unit, responsible for the FFT/IFFT processing, sends to/receives from each core half of the total processed subcarriers.

## 4. Hardware Accelerators

This section describes the different hardwired units highlighting how they scale as a function of the configuration.

### 4.1. Radio Interface Unit

The wireless local area networks per the IEEE 802.11 specifications employ CSMA/CA (carrier sense mechanism with collision avoidance) mechanism for multi-device networks. Each device is supposed to listen to the channel and confirm that there is no on-going transmission prior to attempting its own transmission. As it cannot 'a priori' know the arrival time of the next frame of interest, by default the Wi-Fi system listens to the medium to generate the Clear Channel Assessment (CCA) indications according to the standard requirements and detect signals of interest (for which it must appropriately set the RF and analog gains in about 4us). These extremely time-critical operations are fully handled by the RIU.

At the heart of the RIU is a micro-coded state-machine controlling highly configurable processing blocks. This high level of programmability significantly eases adaptation to a specific RF implementation (specific gain stage split between LNA, down-mixer, VGA) and allows fine optimization on silicon.

On the data-path, the front-end unit is responsible for the Tx/Rx digital up-sampling/down-sampling between the fixed DAC/ADC sampling rate and frequency domain processing rate that varies as a function of the frame bandwidth. It is also responsible for the frequency shift to/from the primary channel. In 80MHz channel operating mode, a fixed ADC/DAC sampling rate of 160MHz is assumed and only the Tx/Rx 20/40/80 filters block is required. If the 160MHz channel operating mode is supported, a Tx/Rx 160 filter is added to adapt to the 320MHz ADC/DAC assumed frequency. The DAC interface can also be easily customized to accommodate higher sampling rates to meet specific radio requirements.

The RIU includes as many front-end units as the number of antennas.

Optionally, the RIU can include a full DSSS/CCK modem which connects directly to the MAC PHY IF without DSP overhead.

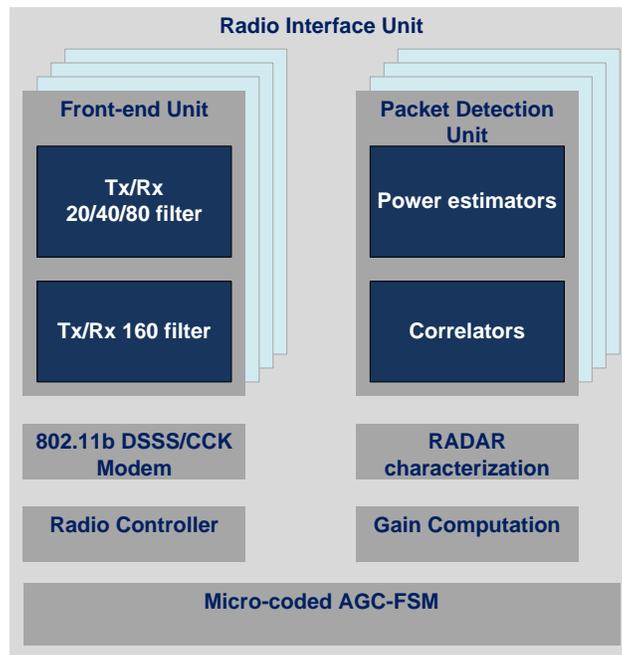


Figure 2 – Radio Interface Unit Architecture

## 4.2. Time Frequency Unit

The Time Frequency Unit (TFU) mainly supplies the FFT/IFFT functionality needed to ease the DSP load.

During reception of the preamble, the TFU simply transfers the data from the RIU to the DSP, which performs all the synchronization estimations. Once the TFU receives the

OFDM symbol boundary estimate from the DSP it can synchronize itself on the OFDM symbol sequence and provide FFT outputs to the DSP.

In 80MHz channel operating mode, FFT size varies between 64, 128 or 256 points as a function of the considered frame bandwidth. If the 160MHz channel operating mode is supported, an additional block that can calculate a 512-point FFT/IFFT from a 256-point FFT/IFFT is required.

On the Rx data path, the TFU can also perform, before the FFT, the time domain DC and frequency offset compensation, offset values to be compensated being provided by the DSP.

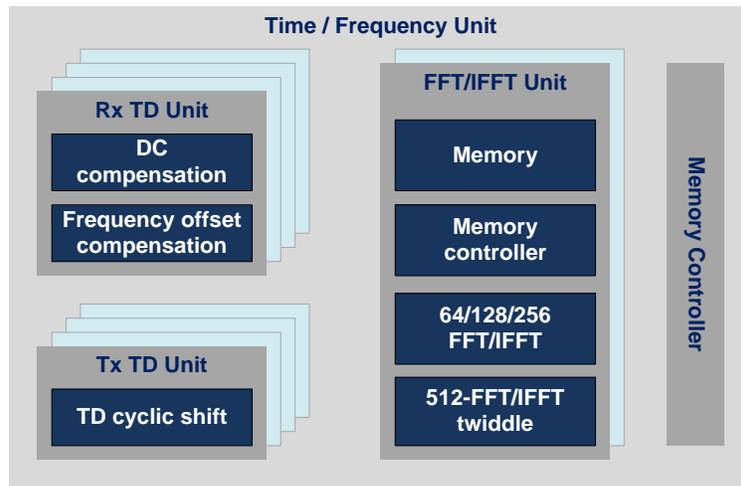


Figure 3 – Time Frequency Unit Architecture

### 4.3. Bit-Processing Unit

The Bit Processing Unit (BPU) performs several bit domain operations, in particular:

- In receive mode it performs part of the de-interleaving, stream multiplexing, convolutional decoding and some of the de-scrambling.
- In transmit mode it performs scrambling, convolution encoding, stream de-multiplexing and some of the interleaving.

The most demanding processing is the convolutional decoding which is handled by several soft-input Viterbi decoders running in parallel. The number of Viterbi decoders is data rate-dependent as defined by the standard.

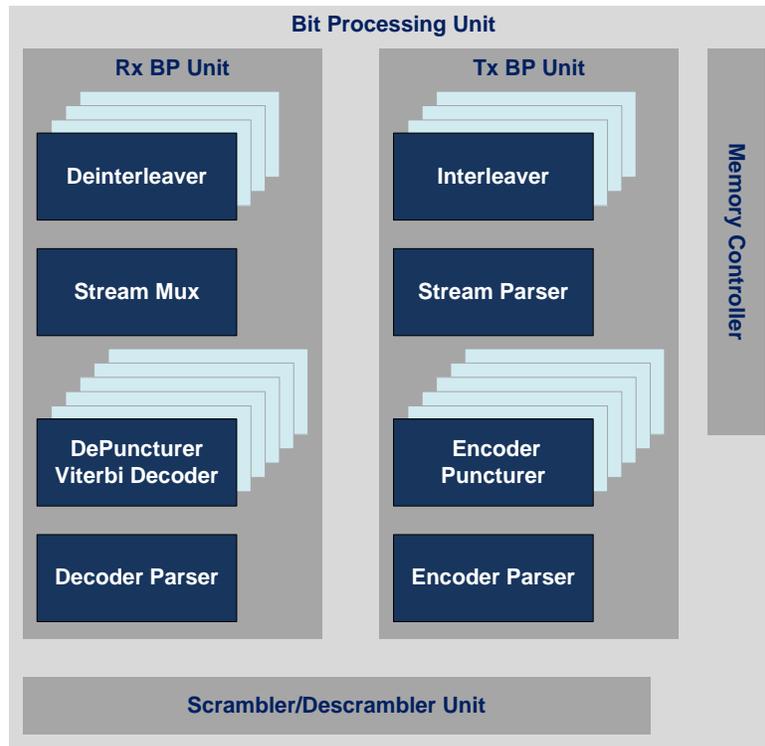


Figure 4 – Bit Processing Unit Architecture

#### 4.4. Smoothing Unit

The Smoothing Unit (SMU) performs the filtering of the channel estimate in the frequency domain to reduce the estimation noise affecting the channel coefficients estimate obtained from the preamble. As a result it significantly improves system sensitivity.

In MIMO operation, each coefficient of the channel matrix is filtered independently (though they can be filtered in parallel). The number of filters instantiated in the SMU depends on the latency requirements. In a configuration requiring several cores running in parallel, there are as many SMU as the number of cores.

#### 4.5. QR decomposition Unit

The QR decomposition Unit (QRU) contributes to the equalizer computation whose complexity corresponds to one of critical path. But, thanks to its generic nature, it is also involved in many other processing tasks including implementation of the SVD of the channel estimate needed to support the beamforming. It is also used in the complex computation of the pre-coding matrix applied as a MU-MIMO AP to handle a MU-MIMO transmission.

As for the SMU, the number of components instantiated depends on the latency requirements, and there are as many QRU as the number of cores.

#### 4.6. MAC-PHY interface Unit

The MacPhy interface Unit (MPU) is responsible for the MAC interface and performs several operations:

- In receive mode, it prepares the Rx-Vector from the SIG fields and provides it to the MAC. It also handles the MacPhy IF and sends the data from the Modem to the MAC.
- In transmit mode, it decodes the Tx-Vector from the MAC and provides contained information to the modem. It prepares the content of the SIG symbols as well as handling the MacPhy IF and sends the data from the MAC to the Modem.

### 5. Example 1: 802.11ac 4x4, 4 Streams, 160MHz Configuration

#### 5.1. Architecture

To realize a 4x4 Wi-Fi MU-MIMO modem supporting 160MHz bandwidth and four receive and transmit streams, the CEVA proposed solution contains:

- One CEVA-XC4210 core
- Radio Interface Unit including four front-end units with 160 filter extension
- Time Frequency Unit with only one 512 FFT/IFFT block
- Bit Processing Unit including six Viterbi decoders
- Channel estimation Smoothing Unit including four instances of the filters
- QR decomposition unit
- Mac/PHY interface unit

The reference software to be implemented on the DSP core is also provided.

This architecture offers a high flexibility margin supporting modification and adaptations of the algorithm to more advanced scenarios and enabling addition of new features i.e. allowing the customer to differentiate their solution.

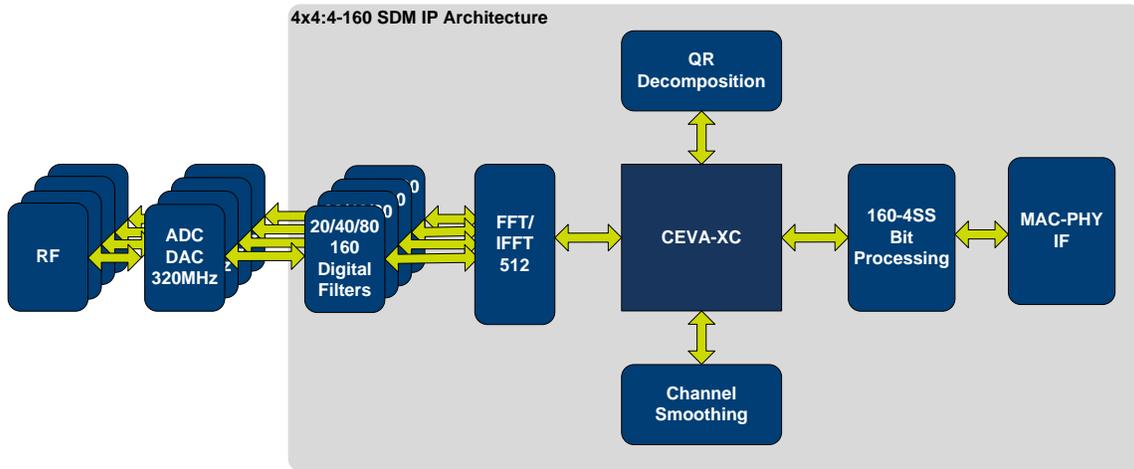


Figure 5 - 4x4:4-160 Architecture Overview

## 5.2 Operations

Based on the previously defined architecture, Figure 6 presents the processing sequence during reception of the VHT portion of a 160MHz 4SS VHT frame.

Each OFDM data symbol output from the RIU is pushed in the memory of TFU hardware FFT, which is launched successively on each antenna. These FFT output samples are then processed by the DSP. The DSP core is specifically responsible for the phase tracking, MIMO equalization and LLR computation. Finally, the soft-bits are provided to the BPU for de-interleaving and decoding.

On the VHT-LTF fields, the DSP computes the 4x4 channel estimates. These 16 channel coefficients are then smoothed over frequency by the SMU. To save time, the process is applied successively on each half of the subcarriers, meaning the first half can be processed by the DSP while those that are pending are smoothed. The smoothed channel is the input of the equalizer coefficient computation algorithms. This complex computation is performed in three steps. First, pre-processing occurs away from the DSP. The QR is then applied and, finally, the post-processing is completed on the DSP.

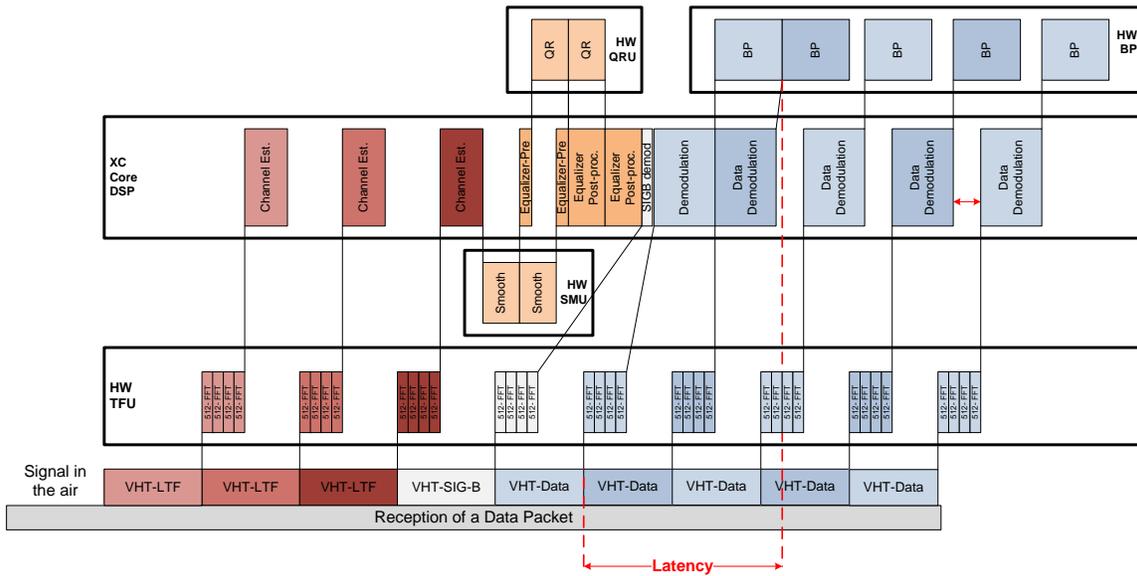


Figure 6 - Timing Diagram

### 5.3 Performance

The timing diagram of Figure 6 illustrates the hardware/software partitioning split during reception and highlights the latency requirements and margin.

It is well known that one of the main constraints of Wi-Fi is linked to the mandatory transmission of controlled frames such as ACK, which takes place after a period of time known as SIFS time after reception of the frame that elicits the response. As depicted schematically in Figure 7 during the 16us of SIFS, the Wi-Fi system must successively complete the demodulation of the receive frame, check the payload at the MAC layer, turn the RF from receive to transmit, all while preparing the next transmission. Since the MAC processing is budgeted at 2us and the RX/TX turnaround at 2us, 12us can be allocated to the PHY Rx processing latency (i.e. the delay between the end of the frame in the air and the decoding of the last bit of the payload).

One of the main contributors to PHY receive latency is the equalizer coefficient computation (the orange boxes of Figure 6) from the VHT-LTF fields. Thus, the latency constraint is particularly stringent for single data symbol frames for which the latency cannot be reduced during reception. Another consequence of the latency constraint is that each system component (DSP & HWA) participating in the data processing must achieve its task in less than 3.6us to prevent the accumulation of latency in excess of 12 us while receiving long frames.

Figure 6 shows that the CEVA Wave 3 Wi-Fi solution offers a large margin compared to the previously defined targets.

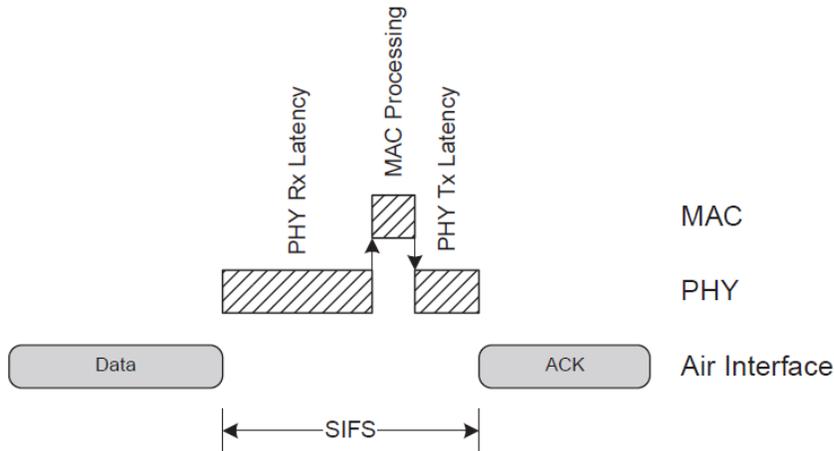


Figure 7 – SIFS Budget

## 6. Example 2: 802.11ac 8x8, 8 streams, 160MHz configuration

The same platform shown in previous section for a 4x4 configuration is expandable to support up to 8x8 with 8 spatial streams thanks to the use of two CEVA-XC4210 DSP cores.

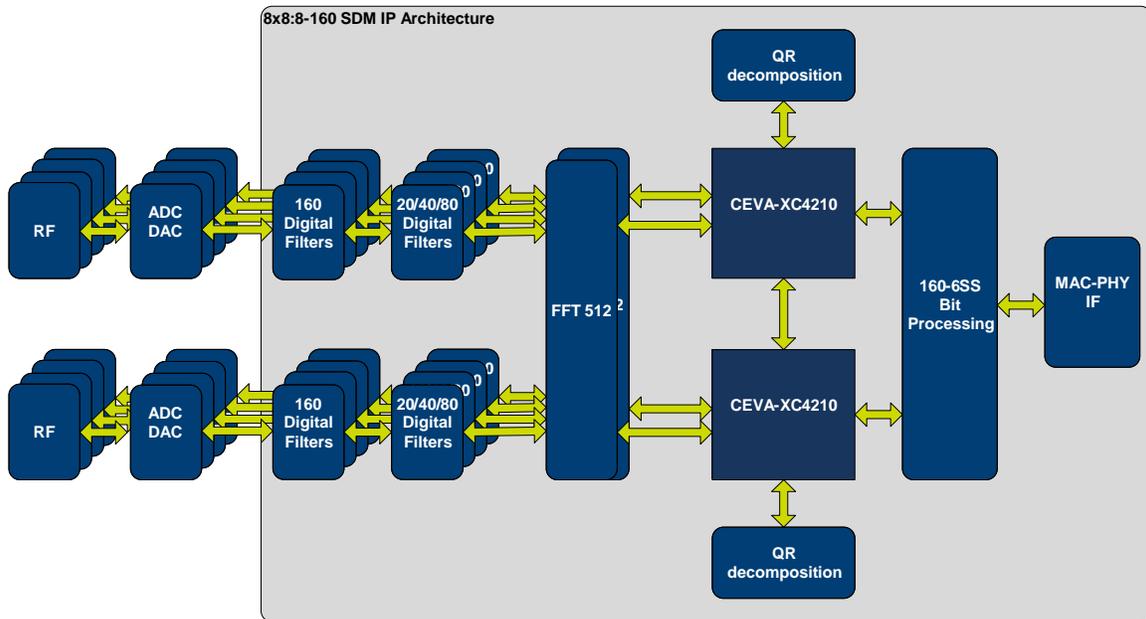


Figure 6 – 8x8-160 architecture Overview

A 10Gb/s 802.11ac Wave 3 solution is achieved thanks to the association of an 8x8 platform with a 4x4 platform to implement an up-to-12 streams dual concurrent 5GHz 802.11ac 8x8 with 2.4GHz 802.11n 4x4.

## 7. Conclusion

This paper presents a flexible and scalable architecture for next generation 802.11 ac Wave 3 implementations that can support complex configurations up to 8x8 MIMO. The single and scalable architecture addresses three different market segments with three different requirements: 2x2, 4x4 and 8x8.

The innovative software-defined architecture, based on a CEVA DSP, provides a lot of flexibility, which is a key differentiator for a MU-MIMO-capable 802.11ac Wi-Fi system. MU-MIMO will soon become an important and an absolute requirement for any 802.11ac device thanks to the higher throughput and better network utilization it brings. There are many different algorithms for precoding and power allocation. MU-MIMO is part of a more complex problem, which is the link-level cross-layer optimization involving MU-MIMO users compatibility evaluation. It includes user selection and scheduling, PER evaluation for Fast Link Adaptation, etc. All of this benefits greatly from a software-defined implementation as more innovative and higher performance algorithms may be implemented over the time.

### References:

1. ABI Research: <https://www.abiresearch.com/press/1391-million-consumer-wi-fi-access-points-shipped/>
2. Machina Research quoted by Qualcomm: <https://www.qualcomm.com/news/ong/2015/01/06/smart-home-happy-family-video>

---

For more information, visit [www.ceva-dsp.com](http://www.ceva-dsp.com) or contact [info@ceva-dsp.com](mailto:info@ceva-dsp.com)