Whitepaper

Advanced Packet Processing for High Performance Access Points

Version	:	1.0
Issue Date	:	20 May 2014
Author	:	Imagination Technologies Limited

Contents

List of Figures

1. Introduction

The data bandwidth requirements of Wi-Fi devices has been increasing 10 fold every 5 years and number of Wi-Fi devices is doubling every year, so the data processing requirement of Wireless Access Points is growing at a compounded rate. With the advent of devices with 802.11ac the trend is expected to continue at a much more rapid pace as it increases data rates by 3x from about 450 Mbps for 3x3 11n to 1.3 Gbps for 11ac.

With the ever increasing data rates and number of devices managed by access points (AP), the solution and the architecture of the AP needs to be scalable, not only with respect to these parameters, but also with respect to the services to be provided at the access points. The Quality of service (QoS) and firewall requirements for Enterprise APs are extremely demanding.

In addition, cloud-based management and controller-lite/controller-less/FAT AP architectures are now common thanks to their elimination of single point of failure, better manageability of the numerous APs etc. This architecture requires the APs to be lot more intelligent and capable in handling data packets by making local decisions.

So, the current generation APs in Enterprise, Out-Door, Small Cell need to be improved for higher throughput, number of stations, security, packet processing, and more. The remainder of this whitepaper tries to present the actual requirements for the APs in these markets and explore ways of providing an optimal solution.

2. Existing Solutions and their Short-Comings

Most of the AP wireless chipsets are developed for client side and retrofitted to APs with a host processor to perform all the functions required in the access point. This approach is not conducive to scaling and is not efficient in terms of power budget, number of security sessions, fast response times, U-APSD/LP latency, and execution of contention free protocols especially with data rates in the 1.3 Gbps range.

2.1. Power Budget

Most deployments of access points require them to be powered by Ethernet (PoE) to avoid additional power connections. Though some of the current versions of PoE provide very high power, most current PoE Ethernet switches (802.2af compliant CAT-5) have a power budget of 14.5W per device. Increasing the number of cores and/or increasing the frequency of the host processor to process the packets exceed the power budgets. In addition, the power consumption of the 802.11ac baseband and RF will increase for the higher data rates.

2.2. Scalability

The number of security sessions supported by the devices is usually 32-64 and this forces the centralized controller (that controls all the APs) to perform the crypto and security functions. A significant portion of the MIPS is consumed in per packet processing (classification, packet editing, and actions like filtering) and with the higher packet/data rates of 802.11ac, the host processor will be a bottleneck.

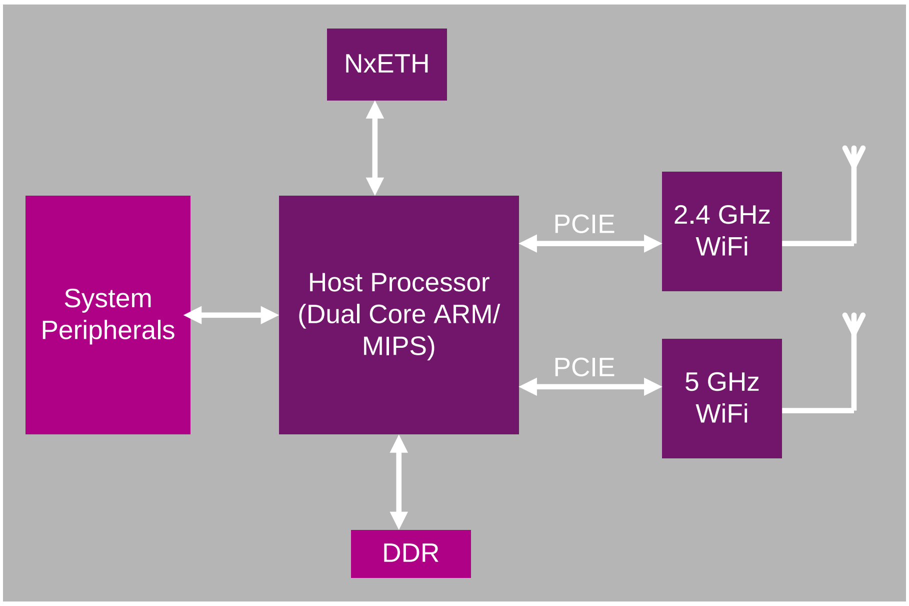

Figure 1: Access Point Hardware Architecture

2.3. QoS

The APs have to support multiple voice, video and data connections. A very sophisticated QoS is required to support these different services across multiple stations/users. Processor based QoS algorithms are known to be sub-optimal in both performance and efficiency in supporting things like voice-video synchronization as they represent two different access categories with large number of clients.

Jitter and sensitivity are important for voice traffic. QoS should be at the edge for it to be effective. Implementing QoS in the controller does not have the same effect as implementing at the edge, as there are a lot of buffering levels in the middle. QoS for short packets is essential mainly to handle voice packets.

2.4. Low Power Client Latency

The client power consumption depends on the response time of the access point. In U-APSD and legacy low power modes, the client wakes up, sends a null data/PS-POLL packet and waits (awake) till it receives a packet. In the AP, the null data/PS-POLL packet is sent to the host and the host processes it and queues the packet to the Wi-Fi client. Most often, the packet for the low power device gets queued behind the already scheduled queues and gets backed up, introducing a lot of latency in the response. Note that most client (same for AP) Wi-Fi chips have only 5-6 hardware queues for 4 ACs and 2 for management/control.

2.5. Spectrum Analysis

As Wi-Fi is in the unlicensed band, the network is uncontrolled and the higher-level protocols need to know about the activity in the current channel and the other adjacent channels. The data collected will be useful for:

- Channel selection
- Scenario replay and diagnosis of packet drops
- Network planning
- Continuous updates to WMM parameters like CW per AC

Some of the current system solutions have an additional Wireless LAN card to address these requirements which is expensive and redundant.

3. Policy Framework in Enterprise, Small Cell and Outdoor APs

Depending on the system vendor there is a distinct policy framework from no-rules to performing deep packet inspection based classification. But across all of them the policies need to be based on

- Mobility policies per user ID
- Applications like SIP, H.323, Skype, MSN
- AP and station location
- Port's subnet
- SSID of the wireless network

It is also essential that the policy framework should be unified across the wired and wireless Layer-2/Layer-3 across the LAN (Ethernet and Wi-Fi) to provide services such as VoWi-Fi and Video over Wi-Fi, which need end-to-end QoS. The Wi-Fi will have additional constraints/policies based on location and the SSID. The markings and mappings would be taken from the DiffServ, 802.1p and extended to WMM (or 802.11e)

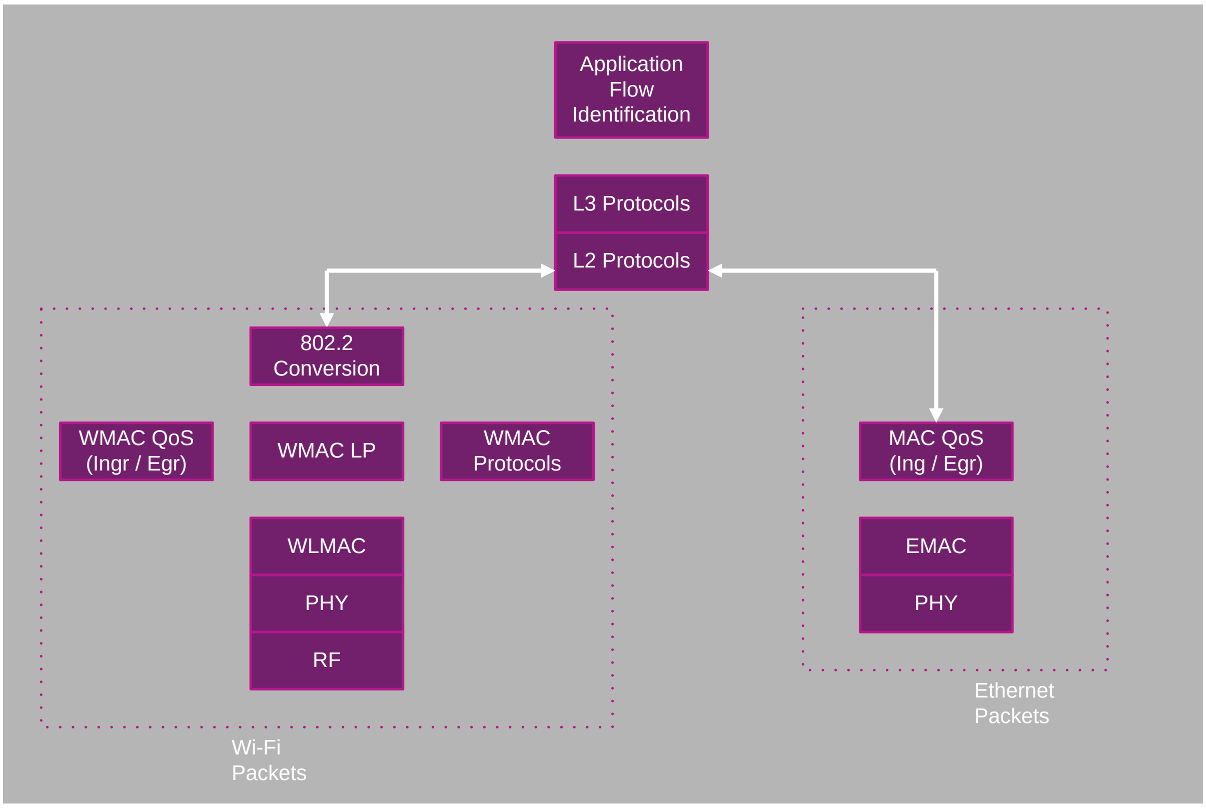

Figure 2: Wired and Wireless Stack Components

However, these policies would still map to a traditional 5 tuple or 6 tuple based on SRC-IP, DST-IP, SPORT, DPORT, PROTOCOL, SSID/VLAN.

The policies can be implemented in a slow path and fast path architecture similar to open flow architecture. The host processor processes all the control packets, connection tracking and initial packets as part of the slow path and sets up the flow action entries for the fast path. The flow action entry has the edit and action fields as mentioned below with appropriate packet headers like MAC header data etc.

Host processor code has protocol helpers that allow connection tracking code to understand protocols, which use multiple network connections (e.g. FTP, H.323, SIP and similar) and mark the ‘child’ connections as being related to the initial connection, usually by reading the related address out of the data stream.

Once the flows are established the packets do not need to go to the host processor for any further processing and they are sent to the Ethernet/WLAN ports directly.

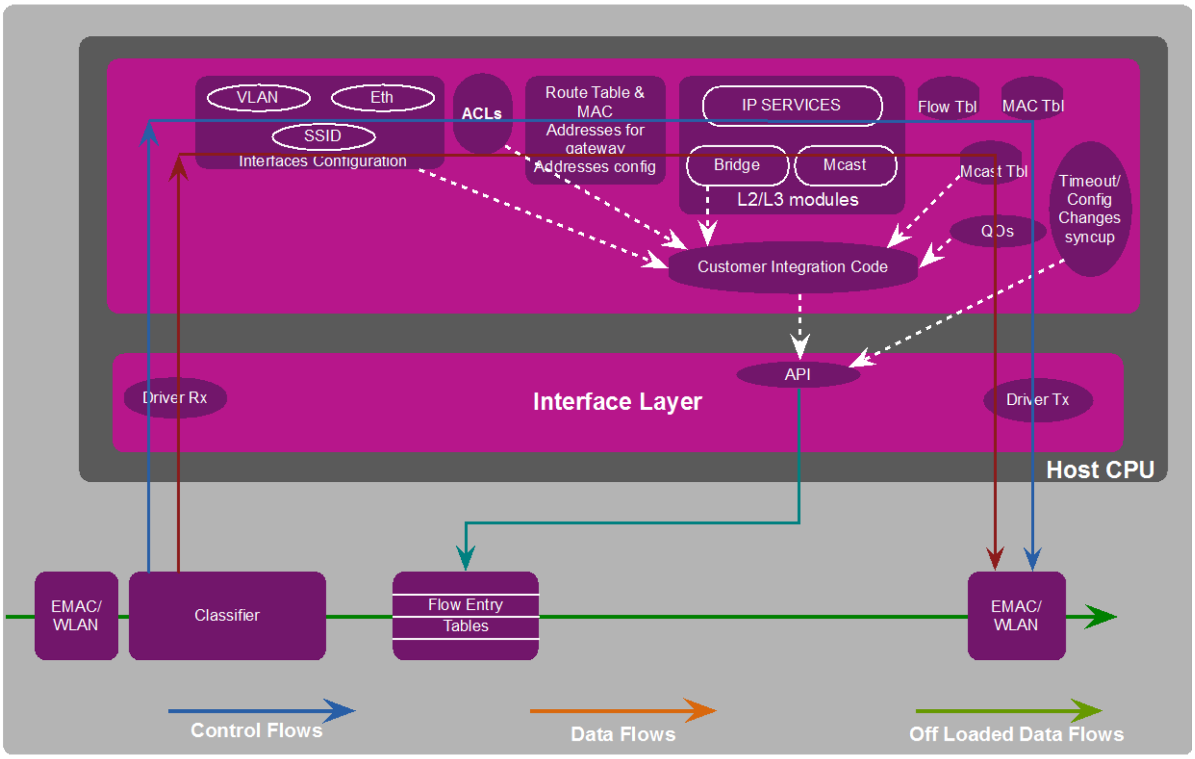

Figure 3: Fast Path / Slow Path Framework

To process the Wi-Fi and Ethernet packets, for implementing the policies, there will be three main stages of processing in the fast path after the initial parsing.

3.1. Classification

Classification parses the incoming packets, extracts the fields like MAC-addresses, source/destination IP address, source/destination ports, protocol, VLAN fields. In addition, this part of the process performs the basic checks for the IP version, TTL etc. A hash value is computed based on the 5/6 tuple and flow action entry is fetched based on the hash value. The flow action entry has the action fields and packet edit options.

3.2. Packet Edit

The packet edit options in the Flow-Action-Entry specify the packet modifications to be performed on the packet.

- 802.11 to 802.3 conversion
- Tunnelling options
- MAC address addition (for route)

3.3. Actions

The actions in the flow action entry field specifies whether to drop, forward or capture the packet. These options are used to realize the firewall and other functionalities.

In addition, the action field also specifies the TOS mapping and QoS rules to apply on the packet. (The slow path of the framework defines the actual mapping, and the action is executed here.)

The packets are then forwarded to the QoS engine, that implement the scheduling and rate shaping algorithms on a per AC per station basis.

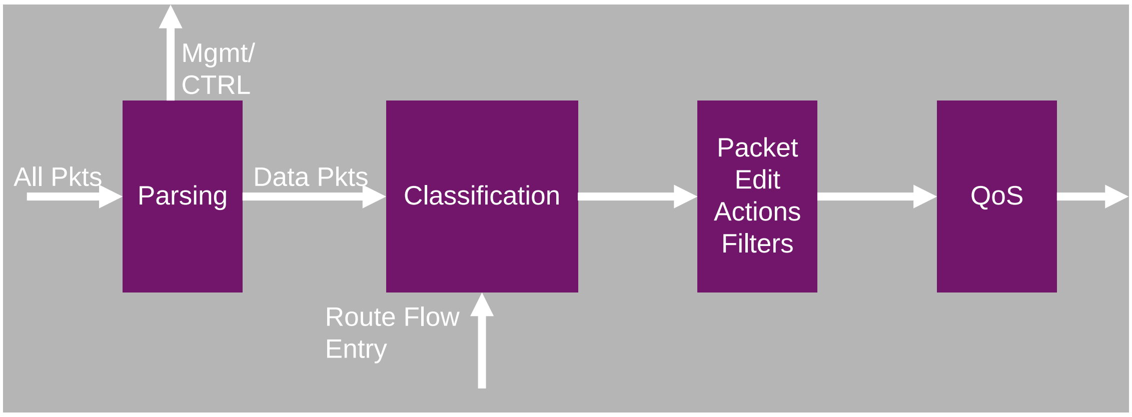

Figure 4: Stages in Policy Framework

4. Proposed Architecture for Future APs

With ever increasing functionality and complexity, dedicated hardware based packet processing in the access points provides much superior performance compared to the software based packet processing. They also offload the host processor from network functions, by implementing the fast path portion of the policy framework shown above, allowing them to serve compute needs.

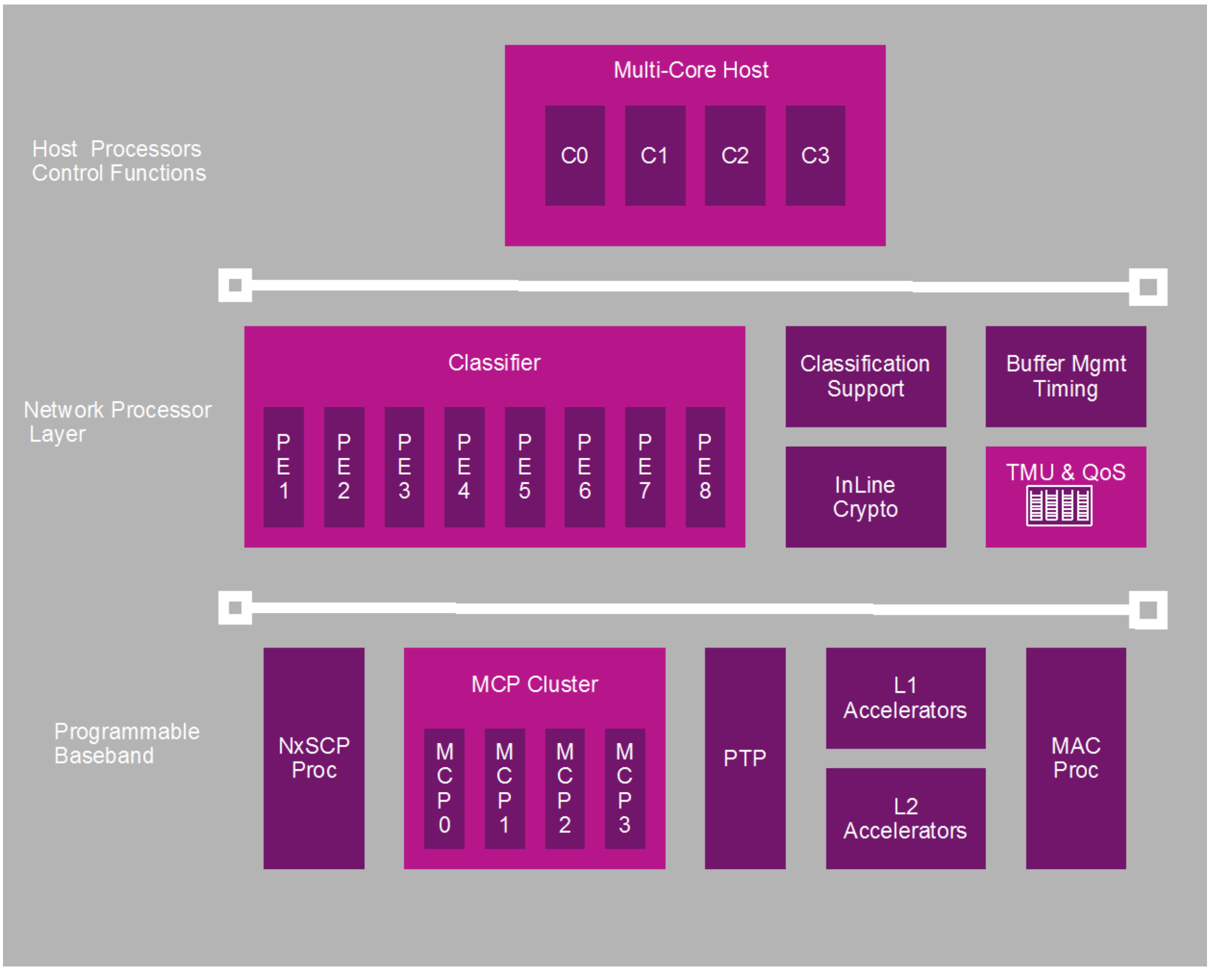

Figure 5: Architecture with Network Processor, Baseband and Host

The following are the different functions that need to be handled differently to existing solutions.

4.1. In-Line Classification

The packets need to be inline classified without making a round trip to the DDR memory. The DDR memory bandwidth is a scarce resource. The classification needs to be programmable owing to the diverse requirements and also changing deployments/features/standards. So a multi-core programmable engine would be best suited for this function.

This would enable the classifier to implement not only static rules/setting, but also dynamic flows as defined in the policy framework.

- DPI
- State full inspection of packets
- Snooping of packets
    - Beyond IGMP to do OSPF etc. for L3 routing
- Application level gateways

4.2. Hardware Assisted QoS

As mentioned previously, processor based or software based QoS is very MIPS intensive and sub-optimal in implementation especially with large number of queues. To provide enterprise class QoS, it is essential to have per AC per STA queue in order not to have one rogue device blocking other STAs and also to restrict the bandwidth across a set of STAs (all STAs in a GUEST SSID) using dedicated hardware. To rate limit at this granularity, a three level hierarchical queuing with rate shapers at each level of the queuing is needed.

- Per AC level
- Per STA level
- Port level

The hierarchical queuing and PerAC/PerSTA queuing also enables video and voice synchronization. The parameters like weights of round robin/fair queuing, shaper token and max counts are dynamically changed as per the rate adaption and the application properties.

Another unique feature of the Wi-Fi, being half duplex in a shared medium, is that the bandwidth allocated to a station could be an aggregate of the traffic in both directions. So when a packet is received from a station (per an AC), the amount of data could be deducted from the rate shaper of that AC+STA. Note that the classification provided above is implemented on both directions.

With this scheme, it is very easy to control and provide guaranteed QoS across a large number of STAs with rules such as

- The total guest traffic will be 1 Mbps
- Priority for certain users
- VoWi-Fi traffic classification and prioritization with shaping

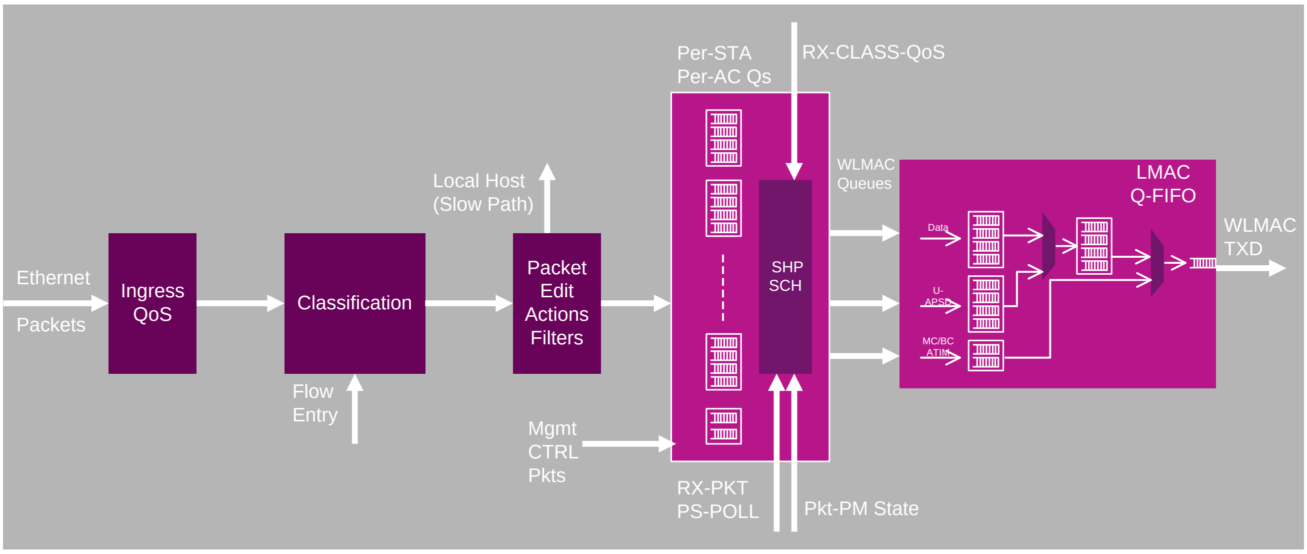

Figure 6: Wi-Fi Transmit Path

4.3. Spectrum Analysis

This feature requires APs to detect and report non-Wi-Fi interferers e.g. Bluetooth, video monitors, cordless phones, microwaves. Typically the expected features include:

- Classify the source of interference (e.g. microwave vs Bluetooth)
- Provide a UI to display real-time signals
- Real-time FFT chart – show energy levels at each frequency component
- FFT duty cycle chart – display the duty cycle of interfering device
- Detect and show if the device is frequency hopping

In addition to the interface from Wi-Fi, the AP needs to detect and report other Wi-Fi BSSs operating on a specified channel or a list of channels and their occupancy/utilization.

During operation on a channel, if this feature is turned on then AP shall detect the interferers without affecting the actual traffic. Alternatively, AP should also be able to initiate a spectral scan on a list of channels and report the results.

Programmable DSP based Baseband would extremely useful for real-time signal analysis to be able to download and run different code for different application environments.

4.4. WiFi IEEE Power Management

The solution has to be very latency sensitive for certain modes of operation like U-APSD. The latency of the access point to respond to a U-APSD trigger, defines the power dissipation of the client as the client initiates PS-POLL packet and waits for the data. The total time to respond should be in the order of 100 µs. Hardware based queues and data transmission right from queues, enables much faster latencies to schedule a packet to the client. In addition special queues are needed in the MAC level is get the U-APSD packets ahead of the normal queues. The TX-QoS block can block/enable queues dynamically based on the incoming packet avoiding the processor to schedule them.

4.5. Packet Coalescing and DDR Store

One of the predominant factors that drives the efficiency of the spectrum/air-time is the ability to burst traffic of a particular station's (and particular TID) traffic using AMSDU and AMPDU. With the help of the hardware based shapers per Q, and by having an additional parameter to increase the burst capability in the shaper, multiple packets from a queue can be scheduled at one scheduling instance.

In addition, owing to the number of stations and the data queues across all the stations, it is preferable that the host processor DDR memory is the only memory used for packet storage. The AMSDU+AMPDU aggregation of 802.11ac provides total packet sizes of up to one megabyte and if multiple buffers are kept internally, the solution would be cost prohibitive. So the architecture needs to traverse through the packet only once and utilize the DDR memory for packet store.

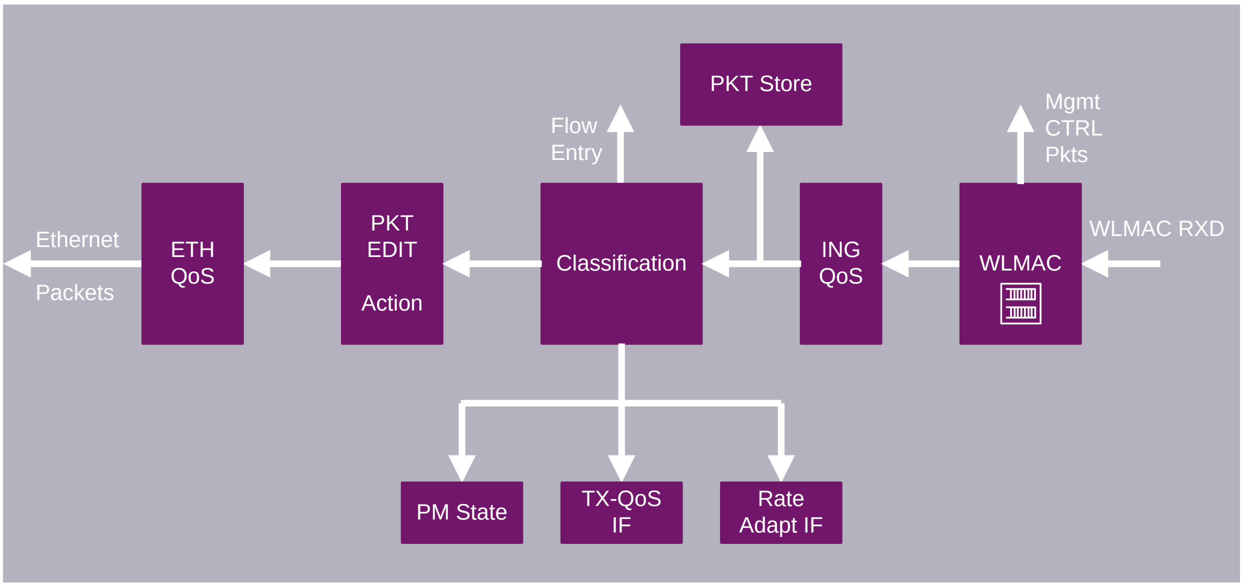

Figure 7: Wi-Fi Receive Path

4.6. Multicast to Unicast Conversion

Wi-Fi as it is specified is not reliable for multicast traffic (no ACKs for multicast). So opportunistically, the access points convert the multicast traffic to unicast to make the communication reliable. Each time this conversion is done there is usually a full packet copy of the data to each of the unicast paths. This is a very expensive operation. A block that can perform the multicast to unicast transfers without having to do individual buffer copies and that can perform header conversion in the datapath is essential to support high rates of multicast traffic (addressed to multiple clients). It is desired to have reference counting in conjunction with this block, instead of having the host processor maintain reference counts per packet.

4.7. Programmability

AP system companies have significant value differentiation in implementing certain protocols like rate adaptation, fast roaming, DFS etc. So the solution should enable reuse of the existing software in these areas. As the deployment scenarios and standards are evolving, it is imperative to have a programmable packet classification and packet editing solution for future upgrades.

Different system solutions could have quite a significant impact in the treatment of very common fields like BSSID treatment etc. so hardware level assumptions to do look ups or classification could prove inadequate. A fine grain software control over the main/fall back rates, transmit power, carrier sense thresholds, contention window parameters is required.

So to accomplish these two tasks, it is desired to have programmability in the MAC layer and also at the packet-processing layer above the lower MAC.

4.8. Scalability

Outdoor, enterprise access points, for example at university campuses, airports and stadiums need to work with hundreds of clients at a time. The AP needs to scale the number of stations with respect to security sessions, QoS across all the stations etc. The 802.11i security on the data packets and protected management packet of 802.11w should be handled in the Wi-Fi chip for scalability and to avoid tunnelling of data packets to the controller. The scalability is required with respect to the number of VLANs, multiple SSIDs and virtualized WLANs.

So, hardware based key look up and context swapping is required to support security sessions up to 256. And to support 64K sessions, 5-6 tuple-based hash is computed and it is used as an index to the flow entry table. Hardware based QoS as described above provides scalability of the traffic management.

4.9. Airtime fairness across stations

Packet schedulers generally perform round robin arbitration across of each of the stations. Stations far away take the same Ethernet bandwidth as stations close by. However, far away stations take lot more airtime. In an AP the difference in data rates of two clients could be 10x (100 Mbps vs 10 Mbps). The airtime usage of the two clients would be 10:1, in turn reducing the overall system performance.

The AP arbitration scheme has to allocate fair time across all the clients. This is accomplished with per station/per AC shapers queues and with rate shapers inline with the rate adapted between the AP and the client. Note that the rate adapted should be the total bandwidth for the two directions as Wi-Fi is half duplex.

5. Changing Dynamics Dictates New Approaches

With ever-increasing functionality and complexity in the access points, dedicated programmable classification, hardware-based queuing and QoS with inter-layer optimization provide superior performance to software-based implementations and also provide tight latency bounds in serving low power clients. The proposed architecture offloads the host processor from most of the per packet functions, across the different OSI layers, allowing them to serve other system specific functions and to be within the PoE budgets.

The Wi-Fi module is the data acquisition point for the Wi-Fi AP and is an optimal place for deploying the functionality described above. It scales with performance needs, eliminates unnecessary data forwarding (tunnelling), and can provide a unified policy framework to implement the much-needed QoS and short latency for low power clients. The proposed architecture is suitable for concurrent dual band in single Wireless LAN module (supported by Imagination Ensigma WiFi Cores) or to support two different Wireless LAN cards for 2.4 GHz and 5 GHz.

The firewall and QoS rules need to be applied on the fly on all the packets at around 1 Gbps, and the firewall rules would be changing based on client mobility etc.

Imagination provides Network Processing Unit IP (NPU), and Radio Processing Unit IP (RPU) to address the packet processing and baseband functions required to provide a wired network experience in a wireless LAN.