Jumbo Frames:
The Microwave Perspective

Abstract

Ethernet frame size is a common discussion item. As opposed to ATM or SDH that have fix size Cells/Frames, Ethernet frames come in different sizes. The frame size variance obviously has an impact on delay, delay variation, and overall performance. In this technical brief, we will focus on non standard Ethernet frames that are used in some applications. These Ethernet frames are sometimes referred to as Jumbo Frames.

This paper will present definitions of Ethernet frame, discuss the reasons for non standard frames and will conclude with best-practice guidelines.

What is an Ethernet frame?

When looking at an Ethernet frame size we usually focus on the following parameters:

1. **Header size** – which may change according to the usage of VLAN tags (4Bytes) 802.1Q or 802.1ad -QinQ. Header size may range between 14 and 22 bytes.
2. **Payload size** - 46-1500 bytes, payload (including data and padding).

To those parameters we need to add the CRC checksum (frame check sequence) of four bytes. Hence, summed up, the size of an Ethernet frame will vary between 64 to 1518 bytes, or as high as 1526 with VLAN tags.

When considering what is the most applicable frame size to be used in the radio, we will refer to 1518 as the common MTU1[1] in the Ethernet world.

For different reasons, Ethernet comes with few overheads per frame. Originally, this was due to the shared media concept. Today it is mainly due to backward compatibly issues.

1. Preamble - 7 bytes

[1] Maximum Transmission Unit

2. Start-of-Frame-Delimiter - 1 byte

3. Interframe gap - 12 bytes

Figure 1 depicts an IEEE 802.3 MAC frame (excluding the VLAN discussion):

Preamble	Start-of-Frame-Delimiter	MAC Destination	MAC Source	Etherty pe Length	Payload (Data and Padding)	CRC32	Interframe Gap
7	1	6	6	2	46-1500	4	12
		64-1518					

Figure 1: Ethernet frame structure

It is noticeable that for small frames, Ethernet adds a significant frame tax. This tax drops fast when arriving at the MTU.

Ethernet overheads also increase the computation requirements, reducing the utilization of some non line rate devices. On the other hand, the larger the MTU size, the higher the jitter and latency. A larger MTU also means higher probability for frame transmission errors, causing a re-transmit of the frame with the assistance of higher protocols, such as TCP/IP.

What are Jumbo Frames and why do we need them?

Jumbo frames are Ethernet frames with more than 1,500 bytes of payload and an MTU exceeding the 1526 bytes defined in the most recent frame extension, IEEE 802.1ad.

An example for further required extensions is enabling common usage of IP/MPLS over Ethernet to deliver Ethernet services. This is referred to as VPLS. Most implementations will require an MTU size of 1540 or 1544 byes to carry the entire overhead, meaning that support for 1600 bytes is a mandatory requirement in Carrier Ethernet networks.

A common Ethernet frame size figure is 1632 for support of Baby Jumbo Frames (1500+{x*4} extension for MPLS labels).In some data applications, system administrators tend to use jumbo frames that can carry up to 9,000 bytes of payload. This is done to improve the utilization of networking and computation resources, as shown in the following table.

	Payload	Overheads	Utilization
64 Bytes	46	38	55%
1518 Bytes	1500	38	97.5%
9000 Bytes	8882	38	99.6%

Table 1: Capacity Utilization of Ethernet frames

It is noticeable from the table above that the larger the frame size, the higher the utilization of the available transmission capacity. Taking into account the reduction of computation overheads (such as TCP segmentation of large chunks of data) the gain from shifting to jumbo frames increases.

On the other hand, it is important to note that the larger the Ethernet frame, the higher the probability of an unrecoverable error -causing re-transmission of the packet in some protocols and thereby reducing the overall capacity. These effects are calculated and shown in Table 2 Below:

	Link Rate (bits)						
Frame Size			10 Mbps	50 Mbps	100 Mbps	200 Mbps	1 Gbps
	Bytes	Bits	10,000,000	50,000,000	100,000,000	200,000,000	1,000,000,000
	64	512	0.0512	0.01024	0.00512	0.00256	0.000512
	1518	12144	1.2144	0.24288	0.12144	0.06072	0.012144
	9000	72000	7.2	1.44	0.72	0.36	0.072

Table 2: Additional delay per single Ethernet frame in different line rates (in Msec)

Table 2 shows that the larger the Ethernet frame, the higher the added delay per frame. When designing a low to mid capacity radio link, network planners need to take the above into consideration and weigh the benefit of capacity utilization against the additional delay variation and probability for errors.

Another perspective for traffic utilization is behavior with TCP protocol. Without going into the details of the mechanism, the protocol increases the capacity until it identifies packet loss. The mechanism then drops back the bandwidth and starts increasing again from half of the capacity. It is assumed that throughput is directly proportional to the Maximum Segment Size (MSS, which is MTU minus TCP/IP headers). This is shown in the calculation below.

Throughput <= ~0.7 * MSS / (rtt * sqrt {packet_loss})

[2]All other things being equal, one can double the throughput by doubling the size of the packet. However this is only true if packet loss is really insignificant, which it never is in real life networks. In fact, in the real world, increasing the size of the packet also increases the probability of losing it. By looking at the above equation it is plain to see that the RTT3 grows proportionally with MSS. Packet loss also grows proportionally with MSS -although this growth is not linear.

The bottom line is that jumbo frames should be used only in very specific scenarios, in which:

1. All network segments are fully controlled and are above 1 Gbps
2. Less sensitive to overall cost
3. Not mixed with real time voice and video traffic

The Microwave Perspective of Jumbo Frames

Microwave links are used for different applications worldwide. However, it is safe to say that nearly 80% of all radio links are used for mobile backhaul. The majority of these links are designed for less than 200 Mbps per link.

One option to increase the throughput in radios is to use intelligent header compression. In radio, every bit over the air counts and, as described in the Ethernet frame definition section, a great deal of capacity is wasted due the header and Interframe gap structure. Modern radios compress Ethernet traffic, thus alleviating overheads almost completely and increasing the effective capacity with short frames.

56MHz/ 256QAM (Bytes)	With Header Compression	W/O header Compression
64 Bytes	530 Mbits	464 Mbits
256 Bytes	403 Mbits	390 Mbits
1512 Bytes	371 Mbits	369 Mbits

Table 3: Effective radio Capacity as a function of the frame size

[2] Round Trip Time

The need for maximizing spectral utilization has brought forth a number of advanced techniques, one of which is ACM (Adaptive Coding & Modulation). ACM ensures that a radio link does not fail in cases of weather-related fading, but rather it will automatically reduce its capacity. This feature helps to ensure the constant and uninterrupted delivery of premium services, but it also makes the microwave system more sensitive to Jumbo Frame effects.

	7 Mhz	10 Mhz	14 Mhz	20 Mhz	28/30 Mhz	56 Mhz
Top Performer PtP Ethernet Radio[3]	>50Mbps	>70Mbps	>110Mbps	>160Mbps	>220 Mbps	>500Mbps

Table 4: Ethernet over licensed microwave spectrum slices: Capacity per Channel

One important aspect to look into is service and traffic patterns in mobile access and aggregation backhaul networks. These imply shorter frames serving voice and control traffic, with clock synchronization constraints, making the use of larger frames even less appealing.

Another perspective is the probability for transmission error. Radio is a reliable media, and through correct implementation, it achieves the highest SLA requirements. Increasing the size of Ethernet frames serves in reality to decrease resource utilization.

To conclude the microwave chapter of this paper let us remember that the MTU was at 1518 bytes as a compromise due to the original low-capacity/high-error-rate network that was available at the time, as well as the available computing resources. While the issue of MTU may be virtually irrelevant in 10 Gbps fiber links, it is still significant in microwave networks – mainly in terms of error rate and capacity, if not in cost of overhead computing.

End-to-End Support Requirements

Modern networks are built globally using different types of devices and media - some self-owned and some over leased services. Introducing Jumbo Frames in the network must also include end-to-end support.

One also has to consider that it is complex to define a designated path for each frame. Thus, when Jumbo Frames are assigned to a certain path, network planners will have to employ

[3] Actual capacity may vary with frame size mix

segmentation and reassembly stations when using very low capacity copper access and even mid capacity microwave links.

Both these issues make the shift to Jumbo frames economically unattractive.

Best Practice Guidelines

- Apply Header Compression: Ethernet microwave, implementing intelligent header compression, achieves significant increase in throughput as compared with the jumbo frames option.
- Fragment Jumbo Frames: When jumbo frames must be transferred through Ethernet microwave, the best option would be to fragment them using a simple router before they reach the indoor unit. Configure the ingress router to fragment the packets (set MTU to 1500), to inter-operate without the advantages of end-to-end jumbo frames.
- Limit the use of Jumbo Frames: Focus on limited scenarios & applications, such as storage, where large transfers are required for backup and data replication.
- Use Standard Frame-size: Try to use standard-sized frames in all Ethernet switches and network interface cards, especially when line speed is less than 1 Gbps (actual speed, not interface speed). If required and relevant, try to stay with the Ethernet frame size that works for MPLS – 1632 bytes (baby jumbo frame - 1500+(x*4) extension for MPLS labels).
- Use Single-ownership Networks: Avoid mobile backhaul in which a mixture of leased and self-owned services are used.
- Protect Premium Services: Avoid scenarios in which mission critical control, voice and video, and other applications are sharing the network.

Summary

Non-standard Jumbo Frames have a high impact on the overall network performance, with a questionable cost improvement. Indeed, Jumbo Frames may increase network utilization in some scenarios with given computational and capacity resources, by reducing the penalties of Ethernet overheads. However, they also introduce a number of significant difficulties that network planners would do best to avoid.

Transporting Jumbo Frames typically results in high jitter in low-to-mid capacity networks. This can impair delay-sensitive applications such as voice and video, and severely reduce the network's overall operation quality. Jumbo Frames also require end-to-end support from all devices and all leased network services, thereby increasing costs.

Although Jumbo Frames have gained some ground in local and storage networks, they are not recommended for broad deployment in mixed environments, and certainly not in general mobile/wireless environments.

About Ceragon Networks

Ceragon Networks Ltd. (NASDAQ: CRNT) is the premier wireless backhaul specialist. Ceragon's high capacity wireless backhaul solutions enable cellular operators and other wireless service providers to deliver 2G/3G and LTE/4G voice and data services that enable smart-phone applications such as Internet browsing, music and video. With unmatched technology and cost innovation, Ceragon's advanced point-to-point microwave systems allow wireless service providers to evolve their networks from circuit-switched and hybrid concepts to all IP networks. Ceragon solutions are designed to support all wireless access technologies, delivering more capacity over longer distances under any given deployment scenario. Ceragon's solutions are deployed by more than 230 service providers of all sizes, and hundreds of private networks in more than 130 countries. Visit Ceragon at www.ceragon.com.