Why Synchronous for Wireless

Why Not Asynchronous?

To understand the benefits of a synchronous protocol, it helps to first look at the disadvantages of an asynchronous protocol. When a node using an asynchronous protocol such as 802.11 wants to transmit a frame, it normally will simply transmit the frame after it senses the channel is idle for a period of time (which is called Carrier Sense Multiple Access, or CSMA). If a collision is determined, due to the lack of an acknowledgment frame, the frame is re-transmitted after waiting an amount of time that increases exponentially for each retransmission. In order to minimize the impact of a collision and to maximize the chance of a successful reception of the data frame, 802.11 includes an optional collision avoidance (CA) function where a short Request-To-Send/Clear-To-Send (RTS/CTS) exchange is first performed, which causes devices overhearing those frames to not access the channel for a period of time. This collision avoidance function may be beneficial in some situations, but it comes with a large overhead and introduces problems of its own, and the impact of these problems is greatly increased in a long-range outdoor system.

Some of the problems associated with carrier sensing (CSMA) and collision avoidance (CA) protocols include:

- **Acknowledgment Overhead**: This is compounded over long distance links due to propagation time.
- **Exponential Back-off**: This is compounded in outdoor networks, where re-transmissions are common due to interference, which causes latency to increase exponentially.
- **"Hidden Nodes"**: This is a classic problem with 802.11 CSMA, where carrier sensing at the transmitter does not sense interference at the receiver. This is greatly compounded in outdoor networks, where obstructions and long distances between the transmitters normally results in them not being able to hear each other.
- **"Exposed Nodes"**: This is a classic problem with 802.11 CA, where the RTS message between a transmitter and receiver causes other potential transmitters to become idle when they could have transmitted successfully to a different receiver. This is greatly compounded in a mesh network, where there are normally many active receivers.
- **CA Overhead**: The collision avoidance overhead due to the RTS-CTS-Data-ACK exchange requires 4 propagation times, which results in large overhead on long-distance links.
- **CSMA Failures**: In a small office or cafe, all stations can normally hear each, which allow them to properly carrier sense and avoid collisions. In an outdoor wireless network, many stations can not normally hear each other, resulting in collisions that cause nodes to experience exponential back-off.
- **Ad-hoc Architecture**: When connecting to an access point in a small office or cafe, all communications occur between the stations and the access point (which is called infrastructure mode) and not directly between stations. This means that most of the transmissions will never collide since all downlink transmissions are from a single device, the access point. In a mesh network using either ad-hoc mode or infrastructure mode there are many simultaneous transmitters and receivers, and all transmissions may collide.
- **Unfairness**: Another classic problem with 802.11 is MAC layer unfairness, and the problem greatly increases in outdoor networks. Due to the increasing back-off during retransmissions, nodes with fewer retransmissions are more likely to gain access to the channel than nodes that are retransmitting. Additionally, nodes that sense the channel becoming idle earlier are more likely to

get access to the channel, and over long distances this results in unfairness to some nodes due to their location.

These problems are basic issues with asynchronous protocols such as 802.11, and all of these problems are drastically increased in outdoor wireless networks. Most people have experienced performance problems related to these issues in offices or cafes, but in outdoor mesh networks the impact of these problems is greatly increased, sometimes resulting in a complete collapse of the MAC layer.

Why Synchronous?

The most obvious reason to choose a synchronous protocol for an outdoor wireless network is to coordinate communications over large coverage areas. Scheduling transmissions not only enhances the efficiency of spectral utilization but also enhances quality of service (QoS) through latency controls, rate control, and traffic prioritization. There are some crude ways to implement scheduled transmissions without being synchronous, such as by simple polling. In fact, 802.11 includes an optional Point Coordination Function (PCF) that uses polling (and 802.11e extends this functionality in its optional Hybrid Coordination Function). Additionally, 802.11 even includes some synchronous features in its base specification, specifically its Time Synchronization Function (TSF), which allows devices to periodically align their clocks that can then be used by functions such as power-save where a sleeping device can periodically wake up at the right moment to see if there is data for it.

However, there are many reasons that 802.11 is not considered a synchronous protocol. Some features traditionally associated with synchronous protocols, such as WiMAX or SkyPilot's SyncMesh™ protocol, include:

- **Contention-less Data Transmissions**: 802.11's base Distributed Coordination Function (DCF) normally puts data in contention, meaning that multiple nodes may transmit simultaneously. WiMAX and SyncMesh schedule data transmissions within time slots, avoiding the contention of data, allowing more bounded latency.
- **Ranging**: DOCSIS (the cable modem standard), WiMAX, and SyncMesh all include a time ranging function, which determines how far apart nodes are in order to compensate for RF propagation at the speed of light. This maximizes efficiency, since inter-frame spaces then do not have to allow for the time of the RF propagation. Synchronous protocols that do not support ranging suffer from this overhead and polling protocols pay the propagation penalty twice. While the speed of light is normally considered fast, on long distance links the 10s of microseconds start to add up, especially as frame transmission times decrease at higher bandwidths and modulations.
- **Periodic Time Slot Grants**: SyncMesh's synchronous nature enables the ability to grant recurring time slots. This means that nodes can be granted extended rights to communicate on certain time slots, which increases efficiency. Asynchronous protocols do not provide this. Periodic time slot grants are useful for providing higher classes of service for applications like Voice over IP (VoIP).
- **Clock Precision**: The features of a synchronous protocol benefit from very precise clocks, which means continually adjusting for phase between time sync messages (or signals from an external clock source) or using very frequent sync messages (SyncMesh performs the former since it is more efficient).

These advanced MAC features are just some of the benefits of using a synchronous protocol, but there is another equally important, if not more important, reason to use a synchronous protocol for broadband wireless mesh – to dynamically point antennas.

Why Synchronous for Wireless

One of the most effective tools an RF engineer uses to improve a wireless link and to minimize a link's impact on others is the use of directional antennas. The benefits of directional antennas include:

- **Increased link budget** (both on transmit and receive), resulting in higher modulation and longer range
- **Decreased interference susceptibility** from external sources
- **Decreased interference** to other systems
- **Increased power** due to point-to-point regulations in many countries

However, the challenge with using directional antennas is just that – they are directional, which requires manual pointing and alignment. In mesh networks, it's advantageous to have 360° omnidirectional coverage. 360° coverage from every node provides easy installation, maximizes redundancy, and avoids expensive and time-consuming system engineering of the mesh.

To provide a node with 360° coverage using directional antennas, multiple antennas are needed, and as the gain of the antennas increases the number of antennas needed to provide 360° coverage also increases. This basic relationship applies no matter what antenna technology is used, from fixed sectors to beam-forming arrays – each of these antenna designs focuses RF energy, and as the antenna gain increases, the RF energy is more focused, decreasing the coverage angle. And while some advanced beam-forming techniques do not use fixed antenna sectors, the RF energy is still focused in a particular direction, so the beam direction needs to be varied in order to provide 360° coverage.

So, most 802.11 mesh networks with directional antennas use manual pointing, where 360° coverage is not provided, and the network must be engineered link-by-link. There has been some research around dynamically pointing antennas with 802.11, but its asynchronous nature prohibits antenna pointing coordination. One challenge with an asynchronous protocol is that some of the transmissions need to be made with omnidirectional antennas (such as omnidirectional Request-To-Send messages), since transmissions are not naturally pre-coordinated. While such a method may allow for higher modulation transmission of the actual data frames, it suffers from decreased range, increased interference, and increased overhead due to the coordination (the latter can be very significant in an outdoor wireless system due to high modulations and the speed-of-light propagation). Alternatively, an asynchronous system could simply use a directional antenna only for transmissions, and use a separate omnidirectional antenna for receptions. The challenge here is that interference is an issue with the receiver, and an omnidirectional receive antenna neither increases the desired signal nor decreases the interference or noise. So, range and link modulation are limited due to the lack of receive antenna gain. Additionally, when only a single side of a link uses a directional antenna, it is not normally classified as a point-to-point link, and many regions limit the effective output power of the link.

By using a fully synchronous protocol, such as SyncMesh, where every communication is coordinated (even bandwidth request opportunities and network entry points), antennas can be pointed on both transmit and receive. This provides all of the benefits of a system consisting entirely of point-to-point links, while still providing the redundancy and simple installation of an omnidirectional system. While these benefits are significant, there are some challenges around creating a fully synchronous mesh protocol.

To summarize so far, there are two primary reasons to use a synchronous, scheduled protocol within a mesh network: MAC layer coordination and to point directional antennas.

Regarding the latter, to avoid the challenges of dynamically pointing antennas, some multi-antenna systems use a separate radio for each antenna (or subset of antennas). This has several problems, with the most obvious problem being cost. Even though there is now the availability of inexpensive 802.11 radios, these radios have many hidden costs due to:

- amplifiers
- increased processing power and processor interconnect
- increased node size
- increased power consumption

However, there is a bigger problem with using multiple radios – self-interference. Even if the radios each use separate frequencies and employ guard bands (which is impractical due to the limited number of channels in many frequency bands), all radios interfere on some level. This can be seen by looking at an 802.11 radio's published adjacent channel rejection values, which is basically the amount of interference from communications on an adjacent non-overlapping channel. The problems due to this self-interference are magnified by the characteristics of outdoor wireless, such as high levels of external interference and weak signal reception due to long links and high amounts of obstruction.

To address the issues of cost and limited channel availability, a reduced number of radios is sometimes used. For instance, some systems use 2 or 3 radios per node. However, a reduced number of radios means a reduced number of antennas, which means either very low gain antennas are used, or 360° coverage is not provided. Both of these restrictions are a large problem for an outdoor mesh system.

To mitigate the interference issues, the most obvious solution is to provide high levels of isolation between the radios and between the antennas. Traditionally, this would mean expensive filters and large amounts of physical shielding which is expensive and increases node size. However, it is impractical to cost effectively provide a sufficient amount of isolation in a mesh node, given typical outdoor wireless scenarios where the received signal may be under -90 dBm while the transmissions might be at +30 dBm. Adjacent, or even alternate, channel rejection along with filters and physical isolation are not enough to provide anywhere near the level of isolation required. So, interference between the radios is not addressed, and results in decreased link modulation and reduction in link range, which are the two main reasons one would use a directional antenna in the first place.

Another general technological issue with using a radio per directional antenna is that such a system can't take advantage of steerable (adaptive beam-forming) antennas. Steerable antenna technology allows an antenna's pattern to be electronically adjusted, so a radio per beam cannot be used since there are no fixed beams.

All of these issues can be addressed by using a synchronous protocol to coordinate all transmissions so that a single radio can be switched among many antennas (or between beam-steering weights). And even though a single radio architecture may not seem to have the capacity of a multiple radio architecture, a multiple radio system cannot take advantage of additional radio capacity due to self-interference. And, the real bottleneck of a mesh network is almost always at the bandwidth injection point (gateway), which means the use of multiple radios in the majority of nodes in a mesh network is wasted money.

Why Not Synchronous?

We've analyzed the benefits of synchronous protocols and the disadvantages of asynchronous protocols in outdoor wireless networks, but what are the disadvantages of using a synchronous protocol? Here are a few disadvantages, and potential solutions:

Why Synchronous for Wireless

- **Clocks need to be synchronized**: Devices participating in a synchronous protocol obviously needed synchronized clocks. This can be provided in several ways, including external clock sources such as GPS or over-the-air clock synchronization. SyncMesh uses a combination of the two, which leverages the accuracy of GPS clocks with the low cost of over-the-air synchronization.
- **Clocks need to be very accurate**: This usually requires expensive clock crystals that are accurate over a wide temperature range. SyncMesh provides an extremely accurate clock source by utilizing an over-the-air calibration protocol along with an internal calibration algorithm that maintains accuracy even with inexpensive crystals.
- **Inefficiencies**: Many synchronous, slotted protocols are inefficient due to their simple Time Division Multiple Access (TDMA) MAC layers, which assigns fixed slots to each user. To overcome this, SyncMesh uses a dynamic slot allocation scheme which assigns all slots in real time.
- **Lack of interoperability with other systems**. Since many outdoor wireless systems leverage unlicensed frequencies, multiple systems may need to share the spectrum. Carrier sensing systems may be able to (in theory) share the spectrum by avoiding simultaneous use, while more complex synchronous systems will not understand each other. However, we've already seen that carrier sensing has issues, and many systems 'tweak' the carrier sensing and back-off protocols to get an unfair advantage over other users of the spectrum. SyncMesh handles multiple users of the spectrum by pointing antennas - the high link budget point-to-point link can avoid interference from other systems, while its directional nature avoids interfering with other systems.
- **Complexity**: WiMAX-like synchronous systems are much more complex that asynchronous 802.11 systems. That is why WiMAX CPEs are more expensive than 802.11 clients, and why WiMAX base stations are significantly more expensive than 802.11 access points. SyncMesh has been developed over a period of 6 years and runs on top of off-the-shelf 802.11 silicon, which lowers cost.

Redefining Broadband Wireless

2055 Laurelwood Road
Santa Clara, CA 95054-2747
408.764.8000
US Toll Free 866 SKYPILOT
sales@skypilot.com