Solving the Wireless Mesh Multi-Hop Dilemma

210-0008-01

Executive Summary

As wireless mesh networks become more popular and their size and complexity continues to grow, mesh networks that contain multiple hops become increasingly vulnerable to problems such as bandwidth degradation, radio interference and network latency. For example, with each hop in the network, throughput can decline by as much as 50% per hop, with degradation quickly compounding over multiple hops, resulting in severe performance degradation across the network. In extreme cases, where voice and video applications are heavily at work, a complete connection loss can occur when latency and RF interference reach unacceptable levels.

A new breed of structured wireless mesh is needed

A traditional mesh network offers limited scalability, and the inherent multi-hop dilemma makes it a questionable candidate for large-scale network deployments. For these larger deployments, a new breed of structured wireless mesh is needed that can provide performance and reliability across the network regardless of the number of hops it employs.

To be effective, the wireless mesh must be a low latency network that dedicates separate wireless bandwidth links to ingress and egress backhaul traffic—similar to a full duplex connection—and utilizes the highest available throughput automatically. It must be a multi-radio, multi-channel and multi-RF system that is both modular and scalable with the flexibility to adapt to any future technology.

This type of low latency mesh network has been tested in “noise-free” (lab) and “real-world” (ambient noise) environments. When stepping the number of hops from 1 to 10, results show that even at 10 hops this network loses only 4% of backhaul throughput under noise-free conditions and just 40% in a real-world environment. Latency test results were equally impressive—well within the acceptable limits for voice (VoIP) and video applications.

The bandwidth degradation test results were compared to the best-case scenario for mesh networks that employ just a single radio for backhaul traffic. In this case, the results show a staggering 80% loss over just 5 hops under noise-free conditions.

StrixSYSTEMS
Networks Without Wires®

Introduction

Wireless LAN (WLAN) or Wi-Fi® has become much more than a solution for small networks. Because of its inherent efficiency and flexibility, wireless has captured the IT mainstream and is now recognized as a viable and effective solution for enterprise and carrier-class networks. Enterprises across all vertical markets are recognizing the tremendous operational benefits of enabling wireless client mobility with data and voice over IP (VoIP) services, often creating specific coverage zones or "hot zones" to better serve their customers. Wireless coverage can also extend to much larger deployments such as metropolitan and regional networks, or even multi-city transportation applications.

Why is Wireless Mesh Necessary?

Wireless mesh focuses on reducing costs and making difficult deployments feasible

Deploying a wireless mesh network offers several advantages over other types of wireless deployments. These advantages are mainly focused on reducing the cost of key factors in any network—installation, maintenance and ongoing operational needs. In some cases, because of the network topology, the lack of wiring infrastructure, or simply the high cost of wiring a customer's indoor or outdoor location, wireless mesh networks are the only feasible way to deploy a network infrastructure.

Historically, there are three Wi-Fi approaches in the industry:

- First generation centralized network model—an unintelligent network with independent access points (APs) wired to the same LAN.
- Second generation centralized model—implemented by the majority of wired switch vendors because it was the easiest extension to their existing switches. The trend with this model was to strip the intelligence from the AP and put it all on the switch; however, this approach introduces many undesirable effects (for example, a single-point-of-failure, bandwidth bottlenecks, and a lack of scalability and flexibility). Furthermore, a new switch must be purchased whenever APs are added and the port limitation of the existing WLAN switch is exceeded. With both of these Wi-Fi approaches, there's one consistent fact—they are not really wireless, just "less" wired. Ethernet to the AP is required.

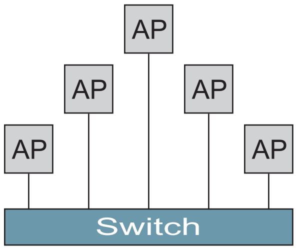

Centralized Network Model

Wi-Fi mesh is designed for optimal performance, reliability, and scalability

- Third generation Wi-Fi mesh networks—an intelligent network where network nodes do not need to be wired to a switch since they can connect wirelessly among themselves via 802.11 links. Mesh network nodes can provide the infrastructure to transport traffic over extended areas or they can provide access for both wireless users and the infrastructure simultaneously. Designed properly, mesh networks are high performance, reliable and redundant, and can be extended to include thousands of devices. This type of network can be installed quickly and doesn't require elaborate planning and site mapping to achieve reliable communications. Simply moving a network node or dropping another node into place can fix a weak signal or dead zone immediately.

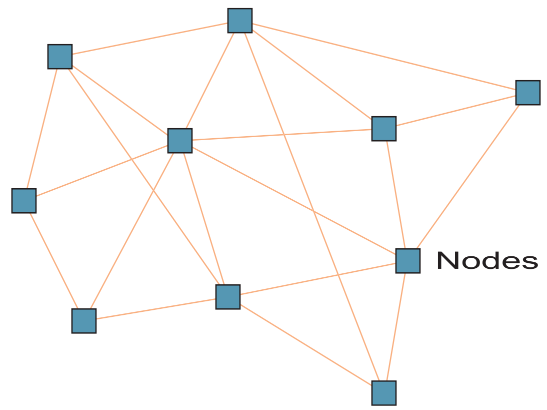

Structured Mesh Network Model

Wi-Fi mesh networks can self-tune and self-heal in real- time

With a wireless mesh, each node in the network establishes the optimal path to its closest neighbor. As the wireless environment changes, such as the addition of a new node or link congestion, data paths are re-evaluated—based on latency, throughput, noise, etc.—and the mesh network self-tunes automatically to maintain its performance at peak levels. If a data path is lost, or if RF interference affects performance, the network self-heals by re-routing traffic so that nodes stay connected and data paths remain optimized. All self-tuning and self-healing processes are dynamic, occurring in the background and in real time—transparent to the user and without the need for human intervention.

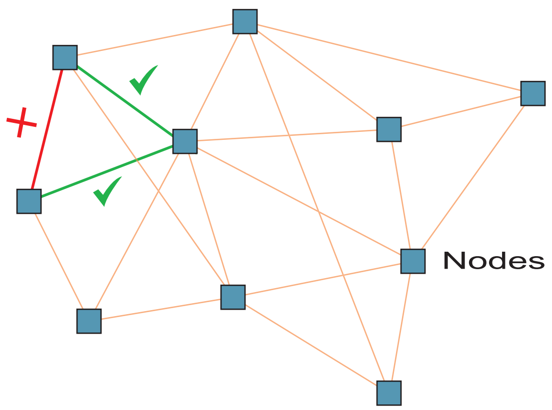

Self-Healing Network

In the outdoors, wireless mesh can avoid obstacles and obstructions

For outdoor deployments, the forwarding capabilities of a mesh architecture allow the wireless network to switch traffic around large physical objects, such as buildings and trees. Instead of attempting to radiate through impeding objects, a wireless mesh network can easily forward packets around an object via intermediate relay nodes. This approach is very useful in dense urban environments that contain many obstructions, or in rural areas where hills or mountains become an obstacle to conventional wireless networks.

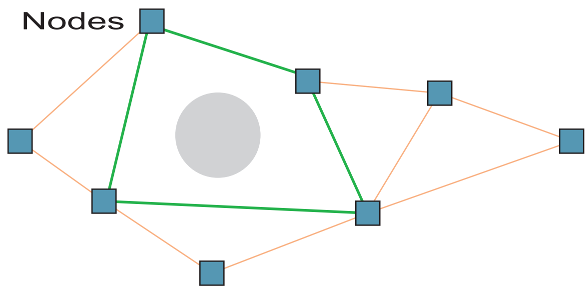

Avoiding Wireless Obstructions

Approaches to Wireless Mesh

Wireless mesh approaches vary, but most have their technology roots in the original concept of the Wireless Distribution System (WDS). WDS is a wireless AP mode that uses wireless bridging, where APs communicate only with each other and don't allow wireless clients to access them, and wireless repeating, where APs communicate with each other and with wireless clients. It is intrinsic to all mesh networks that user traffic must travel through several nodes before exiting the network (for example, the wired LAN). The number of hops that user traffic must make to reach its destination will depend on the network design, the length of the links, the technology used, and other variables.

The single radio approach is similar to a one half-duplex connection shared by all traffic (wireless users, access points, and the wireless backhaul)

The Single Radio Approach—Everything on the Same Channel

The single radio model is the weakest approach to wireless mesh. It uses only one radio on a single channel in the access point, shared by the wireless clients and the backhaul traffic (forwarded between two APs).

As more APs are added to the network, a higher percentage of the radio bandwidth is dedicated to repeating backhaul traffic, leaving very little capacity for wireless clients. This phenomenon is due to the simple fact that wireless is a shared medium. Consequently, an AP cannot send and receive at the same time or send when another AP within range is transmitting. This contention for the available shared bandwidth is based on Ethernet-like collision avoidance rules for wireless (CSMA/CA).

Simple math shows that only limited throughput is possible per wireless client for the single radio approach. For example, if you have 5 APs with only 20 wireless clients connected to each AP, with all APs and clients sharing the same 802.11b channel (5 Mbps), that equates to less than 50 Kbps per user—worse than a dial-up connection. And since all wireless clients and APs must operate on the same channel, network contention and RF interference results in unpredictable latency.

The dual radio approach offers minimal improvement, with the bottleneck still existing in the shared backhaul connection

The Dual Radio Approach—Sharing the Backhaul

With the dual radio approach, one radio is dedicated to wireless client support while the other radio is dedicated to wireless backhaul support—with the backhaul channel being shared for both ingress and egress traffic. What does this mean? A slight improvement is gained for wireless clients, but overall network performance is still poor for the reason that the backhaul is the bottleneck.

The Multi Radio Approach—A Structured Wireless Mesh

In the multi-radio—or structured mesh—approach, there are several dedicated link interfaces of least three radios in use per network node, including one radio for wireless client traffic, a second radio for ingress wireless backhaul traffic, and a third radio for egress backhaul traffic. This approach to wireless mesh networking offers significantly better performance than either the single or dual radio approaches. It allows for dedicated backhaul links that can transmit and receive simultaneously because each link is on a separate channel.

The multi-radio approach eliminates the wireless backhaul bottleneck by providing dedicated wireless links for both ingress and egress traffic

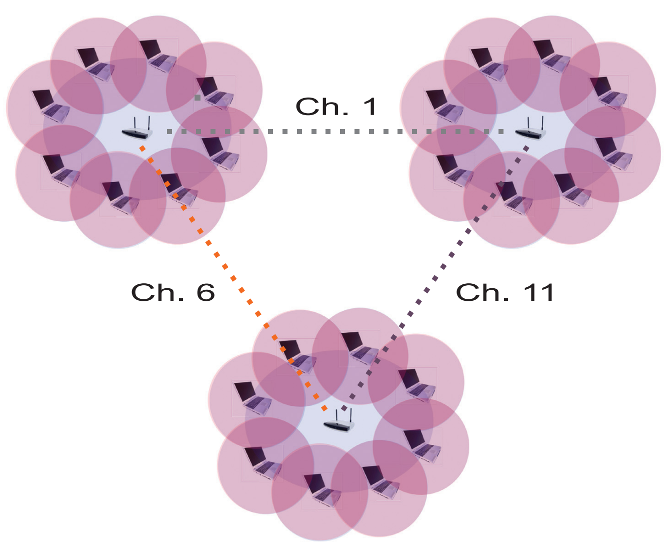

Separate Ingress / Egress Backhaul Links

The Multi-Hop Dilemma

The multi-hop dilemma can be defined as bandwidth degradation, radio interference and network latency problems caused by multiple traffic "hops" within a wireless mesh network.

Bandwidth Degradation

With each hop in a shared backhaul solution, bandwidth degrades by as much as 50%

The problem of multi-hop bandwidth degradation is most severe when the backhaul is shared, as in the single and dual radio approaches. In these cases, each time the aggregated traffic “hops” from AP to AP the throughput is almost cut in half. There are two main theories for this degradation pattern.

Whether you take the best-case theory of **1/n** degradation, where **n** is the number of hops, or the worst-case theory of $\mathbf{1/2^{n-1}}$ degradation, the amount of bandwidth degradation is still substantial, as shown in the following table.

AP Throughput Capacity		Number of Hops									
		1	2	3	4	5	6	7	8	9	10
Best Case	1/n	100%	50%	33%	25%	20%	17%	14%	13%	11%	10%
Worst Case	$(1/2)^{n-1}$	100%	50%	25%	13%	6%	3%	2%	1%	0%	0%

The best-case scenario assumes that all of the nodes are arranged in a linear fashion, similar to a string of pearls, where each node can only communicate with its two neighboring nodes. But in a real-world mesh deployment, any node can “hear” at least three or four neighboring nodes. This is where bandwidth degradation more closely resembles the worst-case scenario. The following chart shows the best-case scenario for single radio throughput degradation for both 802.11a/g and 802.11b.

For 802.11a wireless, the effective throughput drops from 24 Mbps to just 2 Mbps after 10 hops

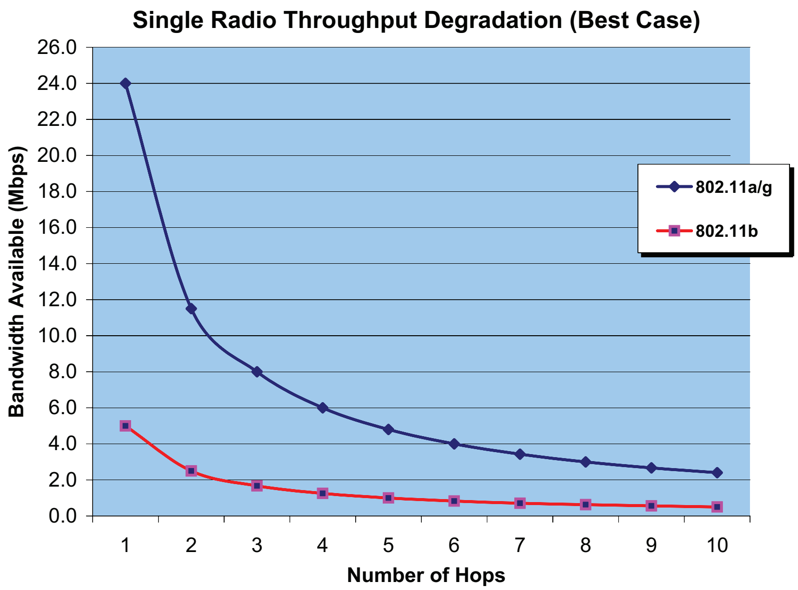

In the case of 802.11b, the starting capacity of 5 Mbps is listed—because in any of its channels, which have raw data rates of 11 Mbps, the effective throughput is closer to 5 Mbps. Similarly, for 802.11a/g the effective throughput is closer to 24 Mbps. As the chart on the previous page clearly shows, even in this best-case scenario the bandwidth loss is unacceptable for medium to large-scale environments.

Radio Interference

The interference from many nearby wireless clients operating in the 802.11b/g spectrum will heavily affect wireless backhaul traffic operating in the same range

Radio interference is a serious issue that affects the performance of any wireless network. It can be defined simply as an undesired signal (or signals) that interfere with the normal operation of other radio communication devices. In today's wireless networks, 802.11b and 802.11g are the two most common technologies used by enterprises and service providers for wireless user coverage. And since a majority of wireless mesh deployments use 802.11b for their wireless backhaul infrastructure, it's easy for their network backhaul bandwidth to become affected by radio interference from neighboring devices that operate in the same band.

Radio interference can also cause transmission errors. These errors can compound. Furthermore, it is important to note that interference may vary in different parts of the network. As previously mentioned, the single and dual radio approaches operate on the same backhaul channel across the network, so when any part of the network is affected by interference, the overall network performance degrades. Also, these approaches cannot accommodate configuration changes (for example, adjusting the channel) in that part of the network that will allow it to avoid the interference.

Wi-Fi networks are not the only radio devices operating in the unlicensed 2.4 GHz and 5 GHz frequency bands. Other types of devices operating in the band include security systems, intercom systems, cordless telephones, and many others. In addition, there are some electronic devices that can leak radio signals in the unlicensed band (for example, microwave ovens, computers, and mobile telephones). These devices can cause interference in a variety of ways. The interference may be temporary (in the case of microwave ovens) or continuous (from a wireless video security camera).

Network Latency

One application that is a key driver for Wi-Fi is voice over IP (VoIP). The creation and deployment of converged Wi-Fi and IP telephony solutions that support both voice and data services is likely to drive the widespread adoption of WLANs throughout the enterprise. In order to support these voice and video applications, there is very little tolerance for network latency and jitter.

For large scale roaming deployments, high network latency and routing protocols can threaten wireless voice and video applications

As a packet traverses the network from node to node, processing delays are naturally introduced. In large wide-area mesh networks, numerous termination points to the wired network would normally be required in order to avoid excessive delay or latency—but this would defeat the benefits of a wireless mesh network. Furthermore, as latency increases, voice and video applications are heavily affected (especially while wireless clients are roaming) which can result in complete connection loss.

Another problem arises when using a Layer 3 (routing) protocol for mesh formation. Using the Layer 3 protocol, where the overhead can be small and data friendly for small to medium deployments, it is not well-suited for large scale environments that support large numbers of users who are roaming and making the most of voice and video applications.

There are various causes of delay or latency that can affect the overall performance of the wireless network. Some of these causes include chatter, congestion, timeouts and retransmissions. Bandwidth Delay Product (BDP) is a concept for measuring the capacity of a network link—something to keep in mind when designing and deploying large networks.

A New Breed of Structured Wireless Mesh

For large-scale wireless network deployments, especially where roaming voice, video and data applications are essential, a new breed of wireless mesh technology is needed—one that offers dedicated wireless links, 802.11a for backhaul traffic, low-latency switching, and cellular-like client coverage. For peak performance at all times these networks must employ a modular, multi-radio, multi-channel, and multi-RF mesh that is very flexible, fully scalable, and ready for any future technologies, such as WiMAX, 802.11n (MIMO), or Ultra-Wideband (UWB).

The ideal wireless mesh network must take a modular approach that contains dedicated backhaul links

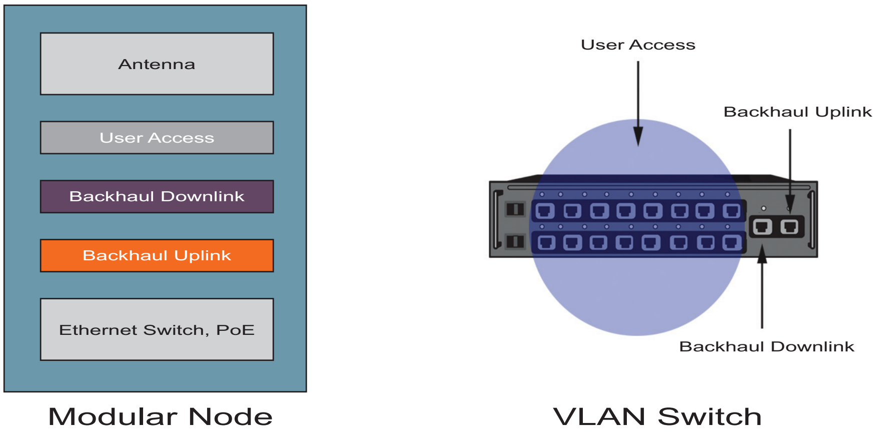

Multi-Interface Wireless Mesh Node vs. Wired Switch

Dedicated Wireless Links

To create a superior mesh network the method used must be different from that used for other typical solutions. It must take a modular approach and have dedicated bandwidth links that use a radio module to perform the function of an AP ("Client Connect") or a Backhaul ("Network Connect"). In the following diagram, note the architectural design similarity between the ideal wireless mesh node and a wired VLAN switch. Both have interfaces for user access and both have dedicated interfaces for backhaul links—a backhaul downlink for Client Connect and a backhaul uplink for Network Connect.

The Client Connect radio module must be designed to receive associations from either the wireless clients (user access) or other mesh nodes (backhaul) in separate modules. The Network Connect radio module must create a link to another mesh node for relaying backhaul traffic (ingress and egress), based on the best possible path to the network egress (the wired world). Consider a repeater node, or a node that extends the reach of your network. This node can consist of three or more radio modules—one for ingress traffic (backhaul), one for egress traffic (backhaul) and another for wireless clients (user access). The following diagram illustrates this concept.

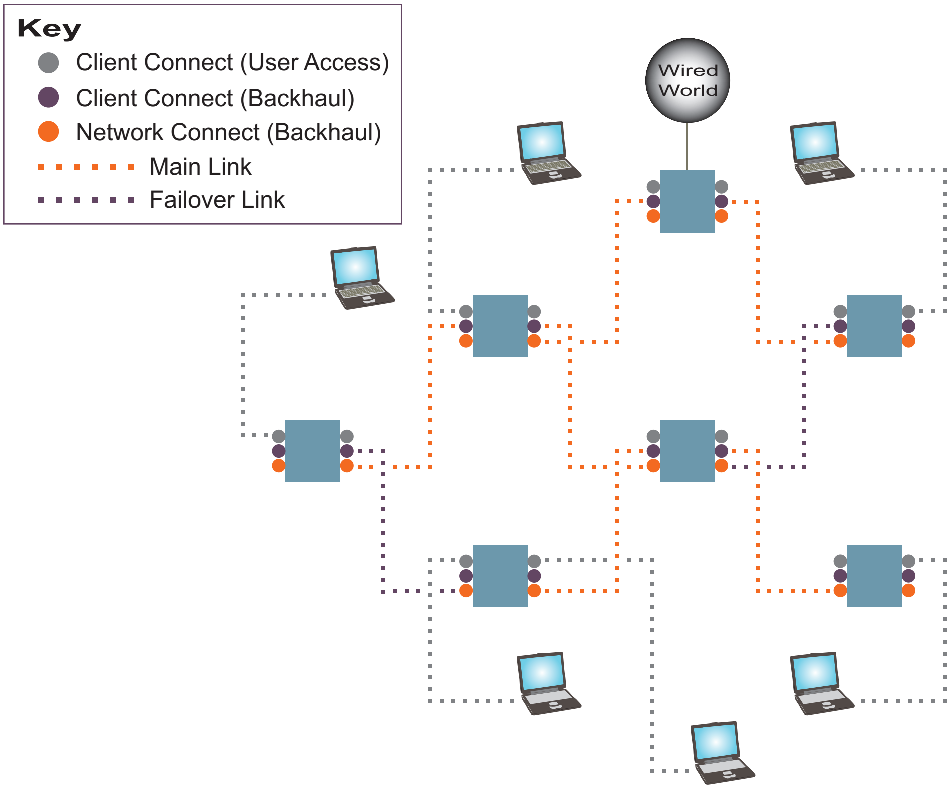

Client Connect and Network Connect

802.11a offers 24 non-overlapping channels whereas 802.11b or 802.11g only provide 3

802.11a Wireless Backhaul

This new breed of structured wireless mesh needs to support different types of wireless technologies. The industry determines what technologies are more widely used for each purpose. As previously mentioned, it is well known that both 802.11b and 802.11g have only 3 possible non-overlapping channels. In contrast, 802.11a does not have this limitation. There is also less radio interference (from other laptops, cordless phones, etc) and a wider spectrum. For these reasons, 802.11a should be used for the backhaul mesh infrastructure formation.

UNII Band	Channel Number	Transmit Frequency
I	36	5.180 GHz
	40	5.200 GHz
	44	5.220 GHz
	48	5.240 GHz
II	52	5.260 GHz
	56	5.280 GHz
	60	5.300 GHz
	64	5.320 GHz
III	100	5.500 GHz
	104	5.520 GHz
	108	5.540 GHz
	112	5.560 GHz
	116	5.580 GHz
	120	5.600 GHz
	124	5.620 GHz
	128	5.640 GHz
	132	5.660 GHz
	136	5.680 GHz
	140	5.700 GHz
IV	149	5.745 GHz
	153	5.765 GHz
	157	5.785 GHz
	161	5.805 GHz
	165	5.825 GHz

The number of permitted 802.11a channels depends on the regulatory domain. For example, the United States FCC regulatory domain reserves four bands in the 5GHz range for unlicensed use. These four bands are designated as Unlicensed National Information Infrastructure (UNII) bands and there are a total of 24 possible channels. Another advantage of 802.11a compared to 802.11b or 802.11g is that all of the available channels are non-overlapping so you can have wireless nodes in adjacent channels and they won't create interference with each other.

In some countries around the world, it may not be possible to take advantage of the 802.11a spectrum. However, in an ideal network the wireless connections between nodes can be made using 802.11g as well as 802.11a. Both provide similar bandwidth and throughput characteristics, and both should support the high-speed modes (Super G for 11g and Turbo A for 11a) to increase your link speed from 54 Mbps to 108 Mbps. The 802.11a Turbo channels are assigned to channels 42, 50, 58, 152, and 160.

A model high speed wireless network must exhibit low latency handoff and sectorized client coverage

Low Latency Switching

The ideal mesh network must accommodate voice applications and fast roaming handoff, which is why the network should be designed around Layer 2 switching. It keeps latency and overhead to a minimum and improves multi-hop performance. Many of today's mesh networks are focused on Layer 3 routing. The problem with this approach is that the routing overhead ultimately limits the scalability of the network.

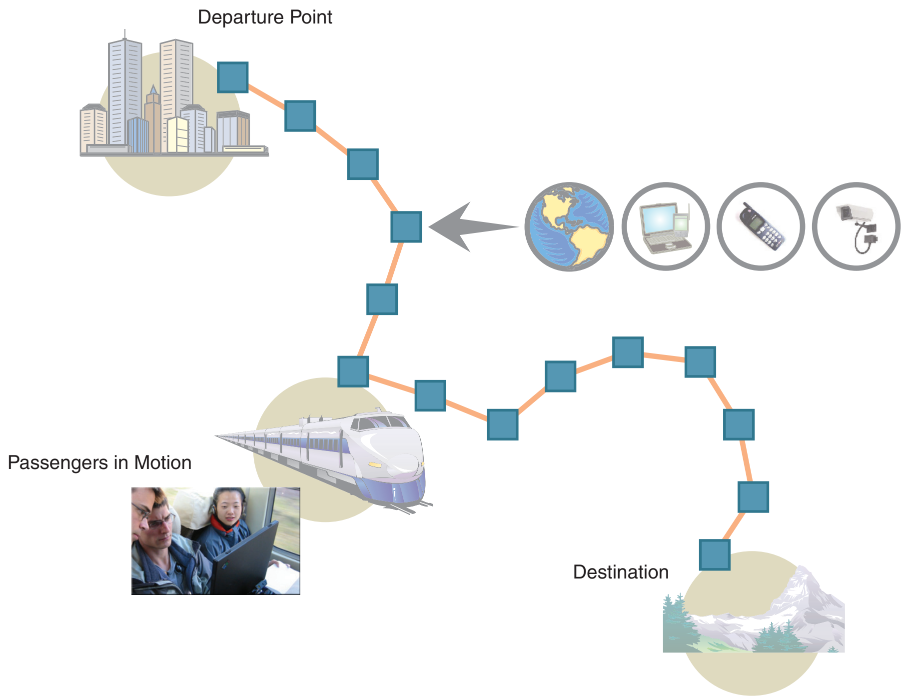

Low Latency Switching in High Speed Applications

Cellular-Like Coverage

The mesh network should also allow the use of multiple, separate "sectorized" antennas (much like cellular radio cell-site antennas) that send signals in different directions, each on different channels, all at the same time. This cellular-like coverage enables simultaneous, collision-free transmission among clients associated with any Client Connect—meaning that more users can associate with the same node at longer ranges and attain a higher overall throughput, because there is less contention with other users.

Proof of Concept

To validate this new breed of wireless mesh network, a series of test scenarios were conducted in both noise-free and real-world environments, stepping the number of hops from 1 to 10. The goal was to characterize traffic degradation through multiple hops that consists of several nodes daisy-chained from a wired node connected to a traffic generation server. This server was used for measuring TCP and UDP bandwidth performance—such as maximum bandwidth—allowing for the tuning of various parameters and characteristics. It also reports bandwidth, delay, jitter, and datagram loss.

Many tests were conducted in both ideal and real-world network environments

Network Test Scenarios

Several types of tests were performed to characterize traffic for all scenarios. These tests are considered accurate when determining what the traffic impact would be when adding or subtracting nodes in an ideal mesh network. The following diagram shows the basic network architecture of this test scenario.

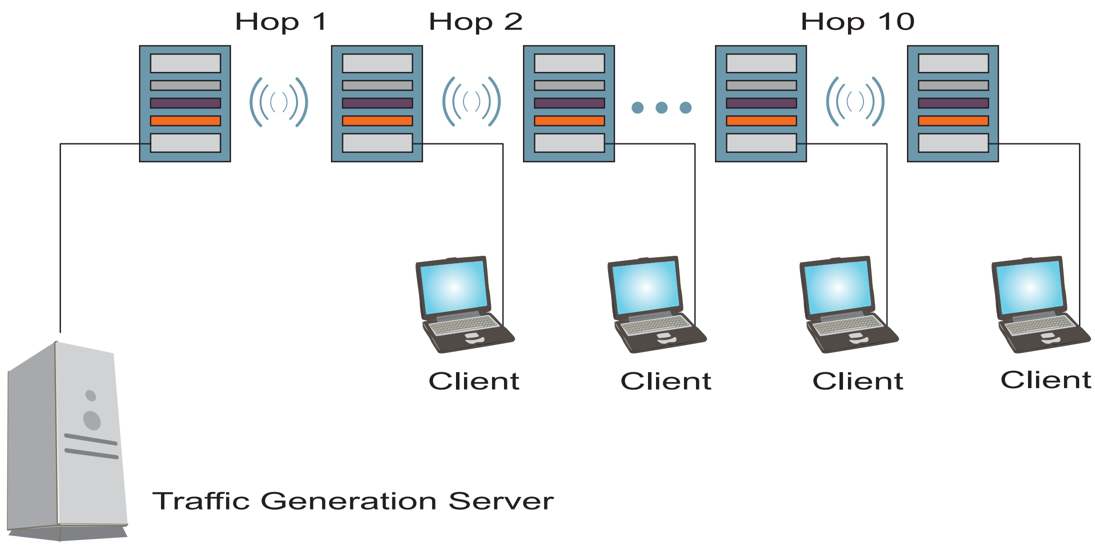

Test Scenarios

The first set of test scenarios were conducted in an environment free of noise and interference, so best-case conditions were assumed. The goal was to accurately analyze the behavior of a mesh network without the influence of external sources or interference. The second set of test scenarios involved real-world conditions, where the network was exposed to the surrounding "noisy" environment. In order to characterize the traffic behavior through a mesh network, all the tests performed used 802.11a for the wireless backhaul.

Network Test Results

Test results were taken at each hop via a wired Ethernet client. Results show that even at 10 hops, this network only loses a total of 4% of its backhaul throughput in a noise-free environment and only 40% in a real-world environment. Latency results are equally impressive with a total delay of only 15 ms with 32-byte packets and 25ms with 1400-byte packets. Both results are well within the acceptable latency limit of 100ms for voice over IP (VoIP) and video applications.

Network Throughput Capacity (Mbps)	Number of Hops									
	1	2	3	4	5	6	7	8	9	10
Noise-Free	23.80	23.80	23.70	23.50	23.40	23.40	23.30	23.20	23.10	22.90
Ambient Noise	21.30	20.15	19.15	18.70	17.80	16.55	15.20	14.25	13.00	12.15

The ideal wireless mesh network must keep wireless backhaul throughput loss to a minimum

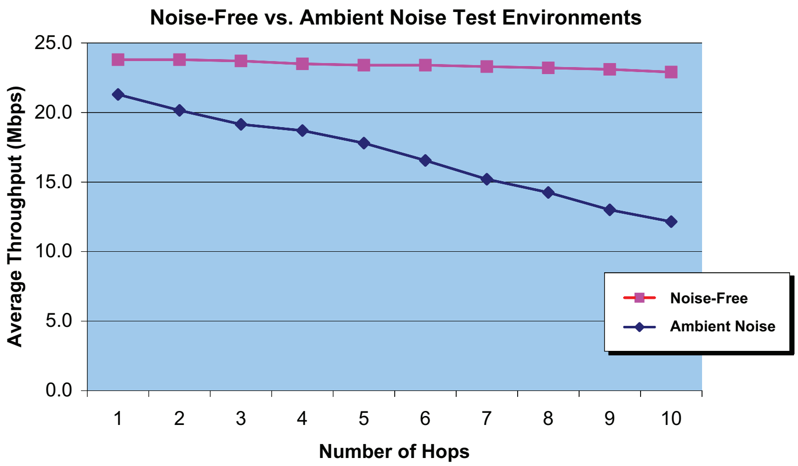

The bandwidth degradation test results were compared to the best-case scenario for single and dual radio mesh networks, where results show a staggering 80% loss over just 5 hops under noise-free conditions.

Conclusion

Large-scale wireless mesh networks with multiple hops are highly susceptible to problems such as bandwidth degradation, radio interference and network latency. If a multi-radio structured wireless mesh is not used, throughput can decline by as much as 50% with every hop a packet traverses. Network latency and roaming are also major concerns, especially in larger deployments. Such latency and bandwidth issues are unacceptable when it comes to wide-area, high-usage environments where wireless clients need to roam (for example, via automobile or train) while using voice and video applications.

The selection of single, dual, and even some multi-radio mesh networks will result in a low performance network that will not scale and which cannot support a wide range of applications.

It is evident that a new breed of wireless mesh network is needed to solve this multi-hop dilemma. This new network needs to employ a modular, radio-agnostic, multi-radio, multi-channel, and multi-RF mesh network system with cellular-like coverage that exhibits superior system backhaul. This will allow for the deployment of high performance, highly scalable networks that support real-time voice, video, and data applications.

About Strix Systems

Strix Systems designs, develops, and markets indoor and outdoor wireless network systems that enhance productivity and efficiency by providing continuous, secure connections to any network, even while roaming with high speed transport (for example, on trains with speeds up to 180 mph).

The Strix Systems Access/One Network family of products represent the next generation of wireless LAN products, leaping past the traditional access point and emerging switch architecture to a true mesh solution—a self-discovering, self-configuring, self-tuning, and self-healing wireless network incorporating robust security and comprehensive management.

This family of products supports all 802.11 RF technologies and offers multi-radio and multi-channel high performance managed mesh networking systems for indoor enterprise deployments and outdoor public safety, municipal and metropolitan networks.

With a modular, radio-agnostic and cellular-like approach, a single Access/One Network node can support up to three user access radio modules in the 2.4GHz or 5GHz bands and up to three 5GHz backhaul radio modules in the 2.4GHz or 5GHz bands that form the multiple point-to-point mesh. In total, each Access/One Network node can support six radio modules on six different channels simultaneously. Furthermore, the network is fully scalable and ready for any future technologies, such as WiMAX, 802.11n (MIMO), or Ultra-Wideband (UWB).

Access/One Networks have been deployed in numerous environments throughout the world, including enterprise, hospitality, education, public safety, housing, and metropolitan area networks. Based in Calabasas, California, the company has developed an extensive portfolio of patent-pending technologies.

For more information about Strix Systems, please visit www.strixsystems.com.