firetide™
instant mesh networks™

TECHNICAL NOTE

Designing a Firetide Instant Mesh Network

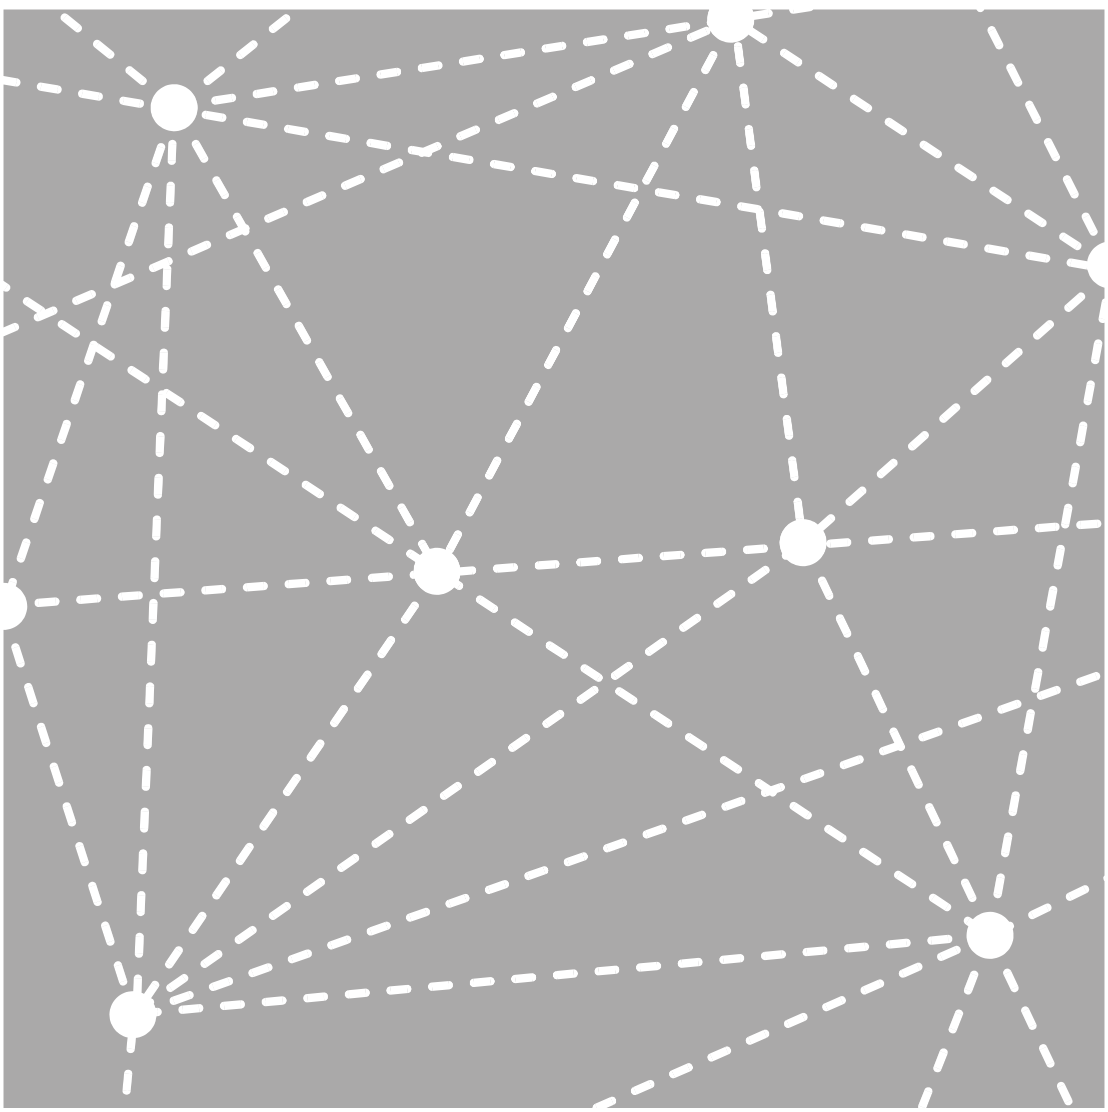

Designing a Firetide Instant Mesh Network

August 2004

Contents

Introduction 3
Key Design Considerations in Wireless Networking 4
General Wireless Considerations 4
Wireless Mesh Network Considerations 5
Designing an Optimal HotPoint Wireless Mesh Network 8
Step 1: Define the Application Requirements 9
Step 2: Create a Preliminary Topology 10
Step 3: Adjust the Topology for the Environment 13
What to Expect During the Implementation 15
Conclusion 16

Firetide, Inc.
www.firetide.com

16795 Lark Avenue, Suite 200
Los Gatos, CA 95032

TN-DAFIMN-083104

Introduction

Designing a wireless network is often more art than science. The science of radio frequency (RF) communications is well understood, of course. But applying that science in a practical way when designing a network can be quite difficult. Understanding the science is also unnecessary with the right design methodology.

The purpose of this Technical Note is to provide some practical guidelines for designing one type of wireless network: the Firetide Instant Mesh Network. To enable a design that is as straightforward as it is solid for Wi-Fi, surveillance and other wireless backhaul applications, the approach taken involves choreographing the art (step-by-step) without delving any deeper than necessary into the science. The result is a reduction of the myriad design factors into the manageable few that normally make the biggest difference.

In many ways, the design of a Firetide Instant Mesh Network is more “forgiving” than it is with other wireless networks. The reason is the self-managing nature of the Firetide mesh. Because the routing protocol used is purpose-built for mesh networking, the mesh becomes both self-configuring and self-healing. As each node is powered up (or relocated), the entire mesh (re)configures itself automatically. And should any node fail or be taken out of service, the mesh heals itself by immediately (and automatically) reconfiguring the topology to take advantage of any available redundant paths. For those who desire additional background information on mesh technology, Firetide offers a separate white paper titled *An Introduction to Wireless Mesh Networking* that is available on the Web at www.firetide.com.

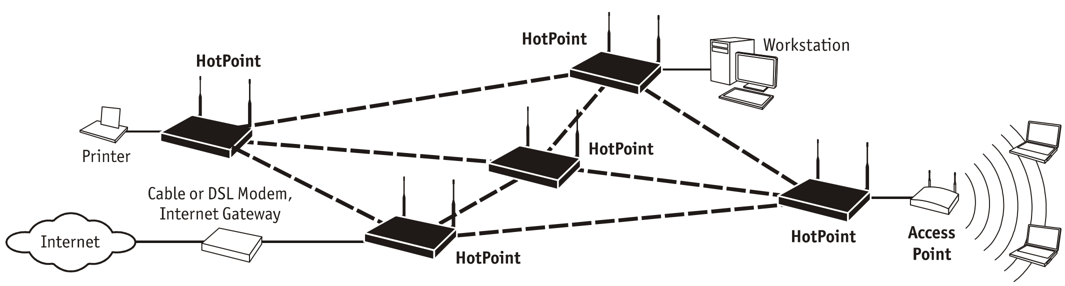

The self-managing nature of a Firetide Instant Mesh Network makes designing the topology a straightforward 3-step process.

Despite the self-managing “plug-and-play” nature of a Firetide Instant Mesh Network, some initial design effort is required to ensure an optimal deployment and to maximize the return on investment. This Technical Note outlines just such a design effort; its content is divided into three additional sections. The first section highlights some key design considerations for wireless networking in general, and wireless meshes in particular. The next section presents a three-step methodology for designing a HotPoint wireless mesh network. A brief conclusion summarizes the main points.

Key Design Considerations in Wireless Networking

This section highlights the key factors that should be taken into account when designing and deploying a wireless network. The considerations are divided into two areas: those that apply to all forms of wireless networking; and those that are specific to wireless mesh networks. While experienced designers may be quite familiar with most of these considerations, some may be new or have a different impact on a mesh topology.

General Wireless Considerations

There are three related factors that are always important in any wireless networking application: radio frequency (RF) interference, environmental conditions and the transmission/reception range of radio-based communications.

RF Interference – The potential for RF interference has increased dramatically over the past few years based on the plethora of all things wireless. This is particularly true in the unlicensed spectrum of 2.4 GHz where most wireless LANs (WLANs), portable phones and many other systems function. For trouble-free operation and peak performance, such interference can fairly easily be minimized or even eliminated based on a judicious allotment of available channels. Ideally no channels will overlap, which totally eliminates interference and yields the best possible performance for both the mesh and the wireless LAN. A channel used by any equipment located "out of range" is, of course, of no consequence. But in radio-rich environments, minimizing channel overlap can require some careful planning.

For example, where Firetide wireless mesh routers are co-located with (or within range of) WLAN access points (APs), maximum channel separation is recommended. Because the entire mesh operates on a single channel, two additional non-overlapping channels are available for the APs. If necessary for spacing of access points, three AP channels can be used, and these should all overlap only minimally or not at all with the mesh channel.

Environmental Conditions – Radio waves in the 2-6 GHz range travel like light waves: in straight lines, whether in a point-to-point or an omni-directional fashion. This unbreakable law of physics requires that antennae, therefore, have an effective "Line of Sight" (LOS) to one another. And as with light, some materials are fairly "transparent" to radio waves; that is, they degrade or attenuate the signal only minimally. Drywall, for example, blocks light entirely but is relatively transparent to radio waves. Denser materials, especially metal or thick concrete walls, block RF transmissions almost entirely. Obstacles that totally block the all-important LOS must be eliminated or circumvented somehow. At a practical level, this requirement might create the need for an intermediate node or nodes to "see around" an immovable obstacle.

A related phenomenon is that radio waves reflect off certain surfaces, which can result in a signal taking multiple paths between nodes. While these reflections can have a significant impact on performance (either negative or positive!), they are very difficult to predict during the design process. At a minimum, however, a prudent designer will be cognizant of such reflections during the implementation stage.

Finally, outdoor networks can encounter additional and variable environmental conditions, ranging from traffic to weather. At radio frequencies below 6 GHz, however, rain and fog have little or no effect.

Radio Range – The energy in a wave diminishes with distance. Some of this effect can be overcome with a more sensitive radio receiver. But background "noise" ultimately makes a weak signal unrecognizable. The implication of this additional law of physics is that systems must be placed "within range" of one another. The range can vary, of course, based on environmental conditions (especially the transparent obstacles) and the particular transmit/receive characteristics of the equipment. But because the law (this time from governments and not physics) regulates the maximum transmit power, there is always an upper limit on the effective range.

Despite these complications, quantifying "within range" is generally fairly straightforward. For example, to operate at the peak IEEE 802.11b "full signal" half-duplex data rate of 11 Mbps, HotPoint nodes should be placed no more than 100-150 M (330-500 feet) apart—assuming, of course, they have a clear or fairly transparent LOS. Greater spacing is possible, but the data rate can drop as a result. With a clearer LOS and special antennae, outdoor equipment can normally be located fairly far apart. For this reason, Firetide's outdoor wireless mesh routers generally achieve high data rates with a spacing of a half to a full mile (800 M to 1.6 kilometers).

An Important Note: The range recommendations made throughout this Technical Note are based on certain real-world assumptions, which make it difficult to reach the *maximum* distances of 1200 M or ¾ mile for the indoor unit, and 3.2 kilometers or 2 miles for the outdoor unit in most applications. In outdoor environments, for example, a relatively clear line of sight is assumed. But trees and crowds of people can attenuate the signal, and the transition from indoors to outdoors always involves a wall, window or roof. Fortunately, glass, drywall and wood siding have only a minimal impact on radio waves. Concrete (especially when freshly poured), and steel or aluminum (including steel-reinforced concrete), however, can have a severe effect on range. In general: the less transparent the obstacle is, the more closely spaced the neighboring nodes need to be to compensate for the signal loss. And never ignore people! People are 70% water, and water (including the moisture in leaves and new concrete) is a lousy medium for radio waves. So be careful to consider the presence of people in that empty convention center, office building or school.

Wireless Mesh Network Considerations

Mesh networks have two unique design considerations in addition to the three general ones just outlined. The fundamental difference between a wireless mesh network and a wireless *local* area network segment is the end-to-end nature of the mesh that normally covers a fairly widespread area, such as:

- Very large indoor spaces, including warehouses, airport terminals and convention centers
- Entire multi-story buildings, including offices, hotels and apartments
- Campus settings, potentially with a combination of indoor and outdoor coverage
- Community-wide environments, including parks, marinas, downtown districts and entire metropolitan areas

As an end-to-end network covering large distances, the two additional design considerations involved in wireless mesh networks include overall performance and degree of redundancy.

Overall Performance – Because wireless mesh routers communicate differently than wireless access points, achieving optimal performance with a wireless mesh involves a different set of variables. The underlying reason (the details of which are beyond the scope of this document) is the use of *Ad Hoc Mode* in the mesh vs. *Infrastructure Mode* for client access. The basic difference (for those who care about these things) is that Infrastructure Mode has clients contending for the access point's bandwidth, which makes a Wi-Fi network contention-based. In an Ad Hoc Mode network, bandwidth is shared across the mesh, much as it is with an Ethernet LAN.

So in addition to the RF environment, the performance of a mesh network is very much dependent on both the IEEE 802.11 standard in general and its Ad Hoc Mode protocol in particular. For starters, because the basic 802.11 protocol is half-duplex, the end-to-end packet throughput is actually *half* of the distance-determined transmission speed achieved. The collision avoidance/recovery mechanism used for Ad Hoc Mode communications (similar to Ethernet's CSMA/CD) also has an impact on throughput as nodes wait to listen before transmitting (to avoid collisions), or recover from actual collisions. The overhead of the Firetide mesh network control protocol is negligible by comparison, consuming less than 2% of available bandwidth. Finally, hop counts of anticipated traffic flows through the mesh may be a consideration with latency-sensitive applications, such as Voice over IP. For this reason, HotPoint routers are designed with a low latency of less than 4 milliseconds, and the mesh protocol used automatically determines the Least Cost Path through the mesh based on this performance consideration.

It is also important to consider the mesh as a multipoint-to-multipoint network when viewed from the outside. Traffic can traverse the mesh between different nodes along different paths. The result of the many possible concurrent traffic flows through the mesh is an "aggregate" bandwidth that often exceeds the maximum data rate of IEEE 802.11 protocol itself! Such is the hallmark of a well-designed mesh topology: many multiple concurrent paths among clients *and* servers, and clients *or* servers (as peers). A software upgrade is planned to help network engineers maximize aggregate bandwidth through intelligent route management and association techniques.

The bottom line: For all of these reasons, a well-designed IEEE 802.11b mesh should deliver between 0.5 and 3.5 Mbps of sustained packet throughput measured along any one path. This level of performance is similar to that of a wireless access network, for example, where individual users experience far less throughput on a sustained basis than the maximum theoretical data rate of the access point. The level of performance achieved across the mesh depends, of course, on a myriad of factors that range from typical traffic flows to the environmental conditions. But anything less than 0.5-3.5 Mbps likely indicates that the mesh needs to be redesigned in whole or in part.

One final note on performance: Security provisions normally have only a negligible impact on performance. Third party virtual private networks (VPNs) and other security provisions have almost no impact on the mesh or its throughput. And while the effect of end-to-end encryption through the mesh depends on the key length, the effect is typically only minor even with strong encryption of 128 bits (with the Advanced Encryption Standard). Which all renders security provisions an unnecessary consideration while designing the mesh topology for optimal performance.

Degree of Redundancy – The resiliency required or desired can affect how "dense" the mesh needs to be; that is, how closely nodes need to be spaced to create the desired number of redundant paths. Spacing mesh routers more closely together increases both the throughput between near-neighbor nodes and increases the number of distant-neighbor nodes, which generally communicate at lower speeds. During operation of the mesh, the routing protocol automatically selects the Least Cost Path through the mesh. But because the mesh is the equivalent of a shared Ethernet-like medium, too much redundancy can adversely affect performance as traffic shares the air. (This important consideration is covered in greater depth in Step 1 below.)

The general recommendation is for every mesh router to have at least *two* neighbors (to eliminate single points of failure) and more as required to achieve additional resiliency or satisfy other design objectives.

Designing an Optimal HotPoint Wireless Mesh Network

The two most important questions involved in designing a wireless mesh network are:

1. How many mesh router nodes are needed to satisfy the requirements of the application(s)?
2. And where should each node be placed to create an optimal mesh topology?

The answer to the first question is simply “enough” to allow neighboring nodes to be spaced “within range” of one another. While this may seem to be a rather flippant answer, the guidelines provided here should yield a fairly accurate number in most situations. The answer to the second question will often be pre-determined or self-evident based on site-specific conditions and constraints. Because a Firetide Instant Mesh Network is self-configuring, node placement becomes the single most critical aspect of designing the mesh.

Fortunately, the self-managing nature of the mesh topology makes it unnecessary to have the perfect answers up-front for either question. The difference between too many nodes (a really “dense” network) and too few nodes (a really “sparse” network) is often no more than a 25-30% variance from the “just right” number (not too dense and not too sparse). The mesh topology automatically reconfigures itself as soon as an unnecessary node is removed or a “missing” node is added. And any node initially placed in a less-than-optimal location can always be reoriented or relocated to better satisfy the application requirements.

With a background knowledge of the basic considerations provided above, designing a Firetide Instant Mesh Network can now be accomplished following a relatively straightforward three-step process:

1. Define the application requirements, including coverage and capacity needed by the mesh network
2. Create a preliminary layout of the mesh topology
3. Evaluate the topography and RF environment of the preliminary layout, and make any necessary changes or adjustments.

Designing for Different Applications

Here are three examples of how different applications lead to different design goals for the mesh topology.

- ***Wi-Fi*** – Wi-Fi HotSpots and HotZones generally handle a significant volume of traffic, but today that traffic is mostly data and is, therefore, bursty by nature. To achieve the maximum data rate, nodes should be spaced no more than 150 M (500 feet) apart in fairly open interior areas, and be much closer (50-100 M or 160-330 feet) where un-reinforced walls exist between nodes. HotZones in the wide open spaces of the great outdoors allow for greater neighboring node distances of up to 400 M (1300 feet) using Firetide's long-range outdoor units. The result in either case is a fairly dense mesh.
- ***Healthcare*** – Healthcare applications demand high data rates (for transmitting digital radiographs, for example) and must accommodate areas with poor propagation characteristics (the rebar-reinforced concrete walls in hospitals). A maximum separation between neighboring nodes may be as small as 50 M (160 feet) for this reason. Multi-story buildings would require at least one node per floor, or more if each floor has a large surface area. As with Wi-Fi applications, the result is a fairly dense mesh.
- ***Campus*** – Covering a campus inevitably involves making the transition from the outdoors to building interiors through walls, windows or roofs. When making each such transitions, use the shorter range of the indoor unit and take into account the building material involved. It often makes sense to mount an outdoor unit on each building with a line of sight to two or more adjacent buildings. This approach makes each indoor/outdoor transition a "given" (thick metal roofs notwithstanding!), and allows the design to focus on connecting the buildings. The result can be a mesh that has both dense and sparse segments.

Step 1: Define the Application Requirements

Every application is a little different, and wireless mesh networks are often used in some rather unique situations. Here is just a partial list of the scenarios where mesh networking is particularly well-suited:

- A wide variety of applications in virtually any industry…
  - Wi-Fi network with multiple WLAN access points
  - "Anywhere Ethernet," perhaps for a video surveillance network
  - Temporary networks that might be needed to handle an emergency situation or facilitate a trade show
- Deployed in these different environments…
  - Large facilities, including warehouses, manufacturing plants and hotels
  - Campuses for corporations, educational institutions, medical centers, government facilities, shipping ports, transportation depots, etc.
  - HotZones, including airports, shopping malls, casinos and convention centers
  - Rural communities and other large outdoor areas, such as resorts, campgrounds, parks, marinas and truck stops
  - HotRegions or other widespread networks that may encompass an entire metropolitan area

Understanding the coverage and capacity required by the application will become a key factor in Step 2. Of particular importance is determining the desired neighboring node spacing, which will establish the transmit/receive radius for placing indoor and/or outdoor HotPoint wireless mesh routers. Different conditions throughout the coverage area may require different spacing to be used; for example, 50 M (160 feet) within a building and up to 400 M (1300 feet) between buildings on the campus. Even within a building, different distances may be needed when communicating through different transparent obstacles. The behavior of Ad Hoc Mode communications (introduced above) is also a factor in this step when the designer must decide between making the mesh dense or sparse—or something in-between.

The Dense Mesh – Intentionally making the network dense is generally a good idea where most traffic is data (e.g. client/server or intranet). Because the nodes are all in fairly close proximity to their next-door neighbors, they communicate at the maximum IEEE 802.11b half-duplex data rate of 11 Mbps. And because data traffic is "bursty" by its nature, the shared medium of Ad Hoc Mode communications operates quite efficiently, just as it does with Ethernet. An additional advantage is that a dense mesh network is also highly resilient.

The Sparse Mesh – Intentionally making the network sparse is generally a good idea where most traffic is voice or video, such as in a surveillance application where transmissions are continuous. The underlying reason is, as before, Ad Hoc Mode's sharing of medium. But in this case the goal is to devote bandwidth as much as possible to unique paths. Or in other words: the goal is to minimize the need to share bandwidth on each hop in the mesh. And with a larger node spacing, there are fewer neighbors to share the air. The individual hops, of course, may operate as low as 1 Mbps owing to the greater distances involved. But that bandwidth is effectively "dedicated" on a sustained basis. Of course, to ensure resiliency of the mesh, each node should still have two neighbors.

Step 2: Create a Preliminary Topology

Start Step 2 by obtaining or creating a *scaled* map or drawing of the topography involved. For a campus or metropolitan setting, the "plot plan" of the area should show all buildings and other potential obstructions. For a multi-story building, a floorplan may be needed for each level. Then, with drawing(s) in hand, start placing nodes.

Begin with the obvious locations first; that is, those determined by the application. These are normally the on/off ramps of the mesh, and may include access points or other aggregation systems at the "edge" of the mesh. The obvious nodes also include any interfaces to servers, LAN switches or WAN internetworking equipment, whether at the "edge" or potentially in the "core" of the mesh. In a multi-story building every floor should have at least one node, and the nodes should be staggered at opposite ends of the building from floor to floor (see Step 3 to understand why).

Next, draw circles in *pencil* (Step 3 again explains why) around all of these nodes. Circles are used because HotPoint mesh routers employ omni-directional antennas to reach as many neighbors as possible, which is a good thing—up to a point. The radius of the circles (drawn to scale) should be *half* of the desired spacing between adjacent nodes. This spacing, which affects the density and is determined by the performance required, should have been established in Step 1. To reiterate: a denser mesh with a higher node-to-node throughput requires a smaller *radius* (25-75 M or 80-250 feet); a larger *radius* (up to 200-400 M or 650-1300 feet) can be used where a sparse mesh is desired. By setting the *radius* at *half* of the desired spacing (which equals the *diameter* of the desired spacing), any circles that touch or overlap (with the nodes meeting "half way) will, therefore, be "within range" of one another.

Now start filling out the mesh by adding nodes for the "gaps" among the on/off ramp nodes. (Note that because these "filler" nodes have no networking equipment attached, their placement requires only a power source.) As before, fill in the fairly obvious gaps or "dead spots" first, drawing a circle around each new node. Because the radius of each circle is half of the desired spacing, the circles should touch or overlap slightly. A useful tip during this step is to place the point of the compass on an existing circle toward the direction of a void. Then strike an arc. Now go to a nearby area and again, with the compass point on an existing circle, strike another arc toward the same direction. A good spot for a node (power availability and other conditions permitting) is somewhere near wherever two or more arcs intersect. Continue to fill out the mesh in this fashion until there are few or no voids remaining.

An Important Note: Designers should be careful at this stage *not* to under-provision the mesh; that is, to make it too sparse even if a sparse mesh is desired. The power and potential of a wireless mesh network derive from the multiple paths available to the routing protocol. After all, it is the mesh topology of the Internet that yields its unprecedented capacity and resiliency. So let the robust routing protocol in a Firetide Instant Mesh Network do its job effectively once the network is up and running. In essence, the ultimate goal of the design effort should really be to give the mesh routing protocol ample paths for aggregating the traffic flows.

Finally, connect the dots. That is, draw lines between neighboring nodes (those where the circles touch or overlap). These lines together indicate the paths available throughout the mesh. The mesh topology and the self-configuring and self-healing nature of a mesh network should be quite evident at this stage.

Additional Node Placement "Rules of Thumb"

Here are some additional guidelines in the form of general "rules of thumb" for placing nodes throughout the mesh topology:

- During the transition from indoors to outdoors (and vice versa) use the desired spacing of the indoor mesh router. The reason is that these neighboring nodes must both transmit and receive, and the indoor units do both over a shorter range.
- The availability of AC power and/or mounting surfaces often determines the exact location of a node. For example, the best possible overlap of circles may find a node located in the middle of a parking lot. But the nearest light pole (with available power) may be 25 meters away. The adjustment may require neighboring nodes to be moved slightly (erased and redrawn on paper) before committing to the design.

- Linear networks (with all nodes in a straight line) offer little or no redundancy, and the aggregate bandwidth is limited to the "weakest link in the chain." So unless the application or environment dictates such a topology, make the mesh two-dimensional by ensuring that every node has at least one alternate path in a different direction. The circles in the Olympic symbol are a good (albeit small) example of such a two-dimensional layout.
- The more nodes the better—up to a point. The reason is that more nodes provide more pathways, which add redundancy and create more possible end-to-end routes through the mesh. Depending on the typical traffic flows of the application, the net result can be an increase in the aggregate performance. A mesh with a greater number of unique neighbors offers better performance with less resiliency. A mesh with more shared neighbors yields better resiliency, potentially at the expense of overall throughput (based on the Ad Hoc Mode protocol used). A "full" mesh, for example, is where every node can communicate with every other node. It is remarkably resilient, but may deliver an unsatisfactory level of performance. For a good compromise between resiliency and performance, each node should have at least two neighbors.
- Finally, remember to consider the total latency of anticipated traffic flows through the mesh. Each router adds about 4 milliseconds of delay, which may or may not be a problem, depending on the application. Total latency can be reduced with more on/off ramps that interface with LANs, WLANs or WANs connecting to clients and servers.

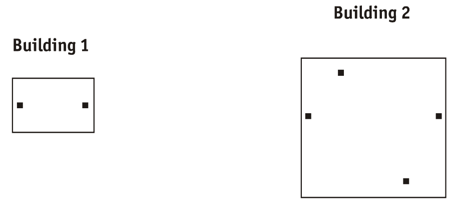

1. Identify the obvious node locations first (on/off ramps, desired access point locations, etc.).

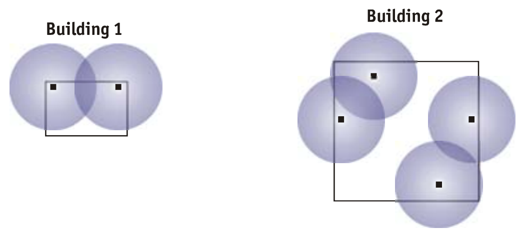

2. Draw circles to indicate wireless coverage of each node (radius = one half of maximum desired spacing between nodes).

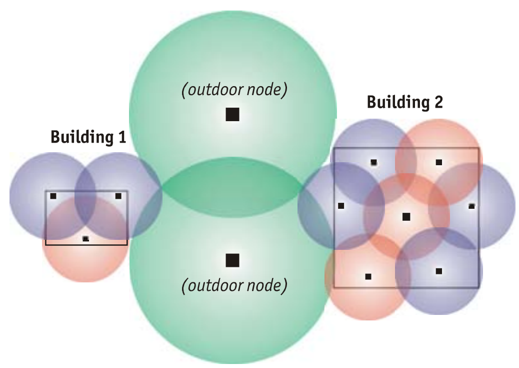

3. Fill in any coverage gaps with additional indoor nodes (shown in red) and outdoor nodes (shown in green).

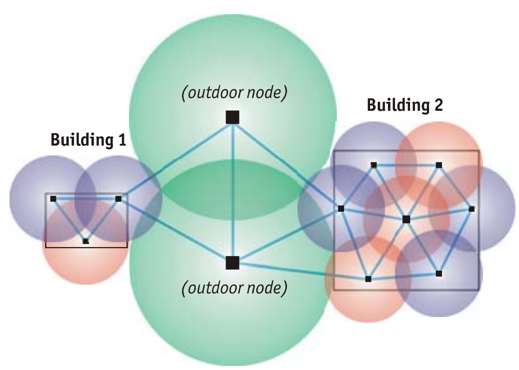

4. Draw lines to connect nodes with intersecting coverage. This will clearly indicate all interconnections among nodes.

Step 3: Adjust the Topology for the Environment

The purpose of this important step is to make any adjustments in the preliminary topology (created in Step 2) that are required to accommodate the behavior of RF communications. Without a site survey, making these adjustments normally involves a bit of guesswork. And even with a detailed survey, the final results may still yield a few surprises! But these guidelines provide proven and practical (and mostly physics-free) ways to complete the design by taking into account two final factors: vertical separation and the RF environment.

Vertical Separation – Begin this step, as before, with the easy part: adding a "third dimension" to your diagram. Circles drawn on a sheet of paper are two-dimensional, which is somewhat desirable because this approximates the transmission pattern of HotPoint antennae. Here are some suggestions for identifying potential adjustments required by the introduction of the third spatial dimension.

Multi-story buildings should require little or no changes in Step 3 if the designer followed the recommendation in Step 2 to stagger nodes from floor to floor. The reason is the mostly horizontal RF transmission pattern. HotPoint nodes are designed with some vertical polarization, which allows the pattern to spread up and down increasingly with distance. Staggering nodes (North side on odd floors and South side on even floors, for example) allows every node to "see" at least two neighbors on the floors above and below.

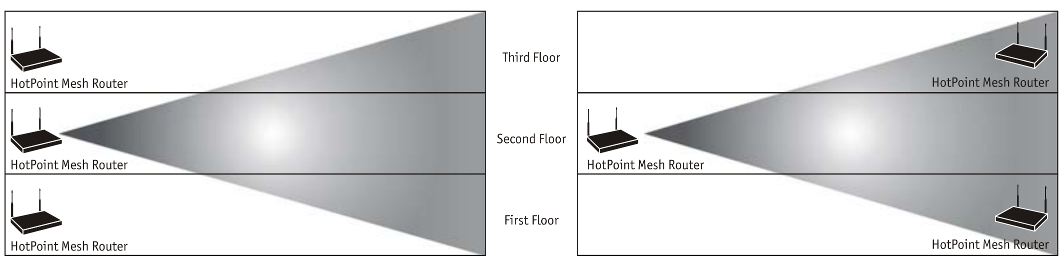

For multi-story applications, staggering the locations of HotPoint mesh routers between floors improves vertical "line-of-sight."

Vertical separations of more than the equivalent of a floor or two, especially over relatively short distances may require a bit more scrutiny. For example, the great outdoors is rarely flat with rolling hills and buildings of varying heights. The result is that every overlapping circle in the preliminary diagram may be at a slightly (or dramatically) different elevation. If neighboring nodes in the preliminary layout are at substantially different elevations (especially if they are fairly close together), there are two potential remedies.

- The first is simply to determine if the outdoor unit (or the antenna on an indoor unit) of one or both neighboring nodes can be angled slightly without making the situation worse for either node's other neighbor(s).

- The second is to add an intermediate node with the antenna oriented to "see" both of its new neighbors. There is a caveat here: The intermediate node (that functions as a "repeater") can create a single point of failure. If this is undesirable or unacceptable, a second "parallel" intermediate node should also be used.

RF Environment – The second part of this step is a little trickier: dealing with RF interference and/or obstacles to RF communications. Interference can normally be minimized through thoughtful channel selection. As mentioned previously, the entire mesh operates on a single channel, and this channel should be relatively clear throughout the topography.

The final (and most difficult) design task involves eliminating or circumventing the RF obstacles, which is mostly a concern with indoor mesh networks. To clear these remaining hurdles, imagine attempting to "see" the neighbors of each and every node in the mesh. (This task is, naturally, only half as hard as it sounds, and like Superman, it is possible to "see" through the RF-transparent objects!) If the view is obstructed by a dense material (thick concrete or metal walls, for example), then an adjustment is probably required. Each obstacle found should be indicated by truncating the circles and/or erasing the line drawn previously. The likely remedy is identical to the one for severe variations in elevation: adding an intermediate node (or a pair of intermediate nodes).

At this point the total number of nodes needed is known, and both of the fundamental questions now have precise answers—at least for the initial deployment. Before finalizing the number of indoor and/or outdoor HotPoint routers required, there is one final consideration: any additional units that might be needed to fill in possible voids in the topology or for use as readily-available spares.

What to Expect During the Implementation

Installing a Firetide Instant Mesh Network is literally as simple as unpacking, mounting and plugging in the HotPoint wireless mesh routers. Nevertheless, a wireless mesh (like most wireless networks) may initially operate somewhat differently than expected. So some "tweaking" may be required before putting the mesh into full production. Fortunately, reconfiguring a Firetide mesh network is as easy as installing one. While not required for the mesh to function, the HotPoint Manager provides a mesh-wide view of the connections among nodes, which is helpful during the tweaking process. HotPoint Manager also provides full centralized control over the administrative and security configuration of the entire mesh.

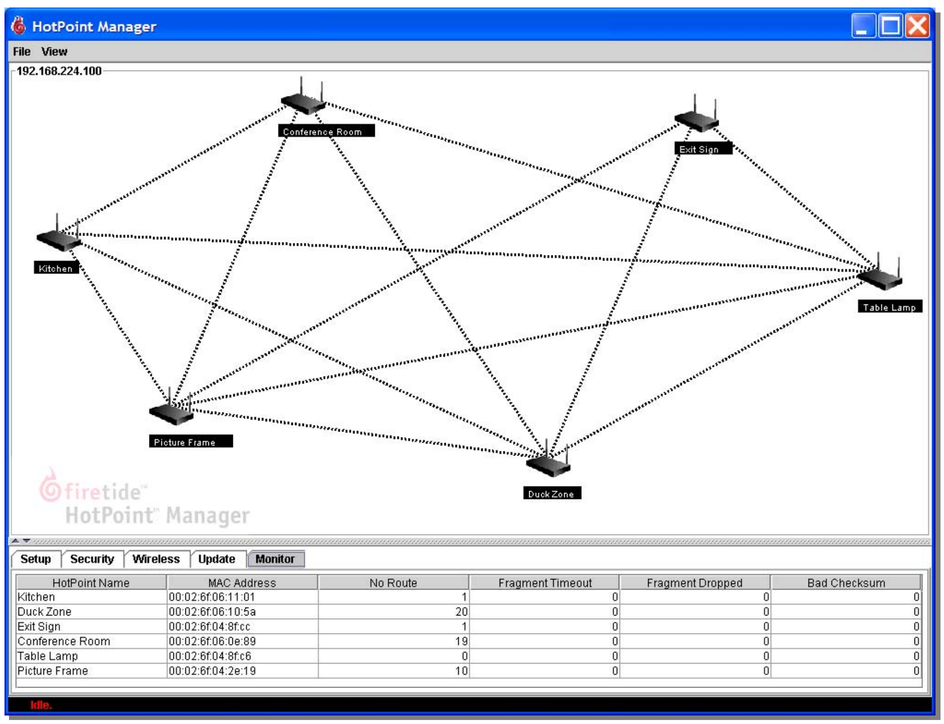

The Hotpoint Manger Software provides a view of the mesh topology that can validate the initial design, and depicts the results of changes, moves and adds during the tweaking process.

The tweaking process involved resembles the familiar add/move/change procedures common to structured wiring for voice and data communications. But with a wireless mesh, these constitute three separate steps and their order is important:

- ***Change*** – Any node's antenna (or the entire enclosure, if necessary) can be repositioned at a slight angle to take better advantage of the two-dimensional omni-directional transmission pattern or to improve the sensitivity to its neighboring node(s). This is the first and easiest remedy to apply because even a minor change can sometimes result in a significant improvement.

- ***Move*** – Any untethered node (those not connected to stationary equipment or Ethernet cabling) can be relocated, provided power is available, to achieve a better line of sight to its neighbor(s). A longer cable or an "extension cord" may be all that is needed to make a major difference in reception.
- ***Add*** – After repositioning antennae and/or relocating existing nodes, new nodes can be added as needed to fill in any remaining voids. Simply power up the new node and observe what happens with the mesh. If the desired results are not achieved, begin this process again by changing the orientation of the antenna or moving the node.

Conclusion

Here are three final suggestions for completing a successful design of a Firetide Instant Mesh Network.

First and foremost is that there is no reason to dread doing the design. A self-managing Firetide mesh is very forgiving and even a sub-optimal design may perform well enough to tweak it using the change/move/add process.

Second, pay particular attention to determining the optimal inter-node spacing. The distance(s) used determine both the mesh performance and its density, which in turn determines the degree of redundancy. A common mistake is to set the desired spacing too high, resulting in a mesh that delivers unsatisfactory levels of throughput and/or resiliency. So when in doubt, be conservative. After all, this is wireless!

Third and finally, seek help. Firetide Hot Fusion Partners are willing and able to help with the design and deployment of any Firetide Instant Mesh Network in any application, anywhere in the world. So call your Hot Fusion Partner or e-mail Firetide at partners@firetide.com to find the right partner for your needs.