First, Second and Third Generation Mesh Architectures

Abstract

Evolving from ad hoc 802.11 networking, earlier generations of wireless mesh provided basic networking over extended outdoor areas. With the emergence of demanding data applications along with video and voice, single-radio "First Generation" single-radio wireless mesh solutions are proving unsatisfactory in many of these demanding environments. Third Generation wireless mesh solutions are based on multi-radio backhauls and deliver 50-1000X better performance, but some custom hardware-oriented approaches limit flexibility and create deployment challenges. Software-oriented Third Generation wireless mesh based on distributed dynamic radio intelligence delivers the same high performance but with the additional benefits of easier installation, better avoidance of interference, and the added flexibility of easy mobility. These new capabilities are enabling many new types of applications beyond the traditional wireless mesh metro/muni environment.

Introduction

Mesh network requirements have evolved from their military origins as requirements have moved from the battlefield to the service provider, and residential networking environments. Today, to cover large areas with a single wired Internet link, more cost effective and efficient means of bandwidth distribution are needed. This implies more relay nodes (hops) than were needed before. Further, growing demands for Video and Voice-over-IP require packets to be moved over the mesh at high speeds with both low latency and low jitter. These new mesh requirements (more hops to cover large areas, more efficient bandwidth distribution and better latency and jitter for Video and VOIP) has given rise to the third-generation of mesh architectures.

Three Generations of Mesh Architectures

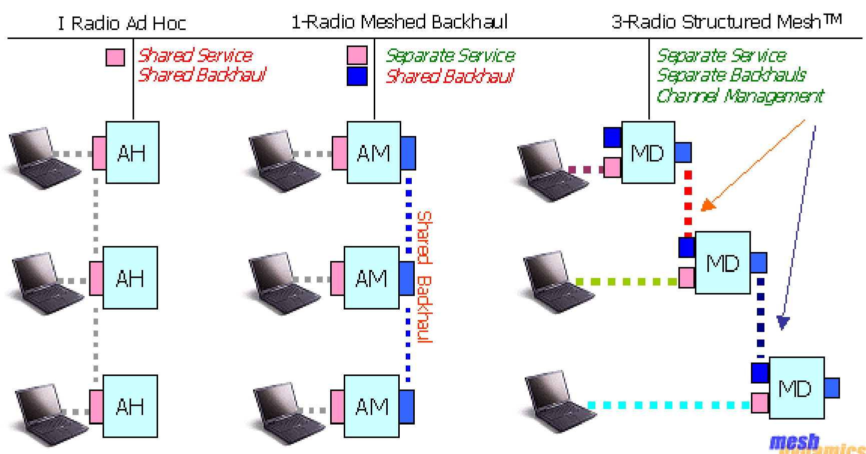

Figure. 1: (L-R): Ad Hoc, 1-Radio Meshed Backhaul, 3-radio Structured Mesh

Three generations of evolving mesh architectures are depicted above. They are (Left to Right):

- First Generation: 1-Radio Ad Hoc Mesh (left). This network uses one radio channel both to service clients and to provide the mesh backhaul. The ad hoc mesh radio, marked AH, provides both services – client access and backhaul. This architecture provides the worst services of all the options, as expected, since both backhaul and service compete for bandwidth.
- Second Generation: Dual-Radio with Single Radio Ad-Hoc meshed backhaul (center). This configuration can also be referred to as a “1+1” network, since each node contains two radios, one to provide service to the clients, and one to create the mesh network for backhaul. The “1+1” appellation indicates that these radios are separate from each other – the radio providing service does not participate in the backhaul, and the radio participating in the backhaul does not provide service to the clients. These two radios can operate in different bands. For example, a 2.4 GHz IEEE 802.11 b/g radio can be used for service and an IEEE 802.11a (5.8 GHz) radio can be used exclusively for backhaul. Though this configuration is sometimes called a Dual Radio Mesh, only one radio participates in the mesh. Performance analysis indicates that separating the service from the backhaul improves performance when compared with conventional ad hoc mesh networks. But since a single radio ad hoc mesh is still servicing the backhaul, packets traveling toward the Internet share bandwidth at each hop along the backhaul path with other interfering mesh backhaul nodes - all-operating on the same channel. This leads to throughput degradations which are not as severe as for the ad-hoc mesh, but which are sizeable nevertheless.
- Third Generation: 3-Radio Structured Mesh (right). The last architecture shown is one that provides separate backhaul and service functionality and dynamically manages channels of all of the radios so that all radios are on non-interfering channels. Performance analysis indicates that this provides the best performance of any of the methods considered here. Note that the two backhaul radios for the 3-radio configuration shown in Figure 1 are of the same type - not to be confused with 1+1 so-called dual radio meshes where one radio is for backhaul) and the other for service. In the 3-radio configuration, 2 radios are providing the up link and down link backhaul functionality, and the third radio is providing service to the clients.

Bandwidth degradation on Single Channel Backhauls

With one backhaul radio available for relaying packets, all nodes communicate with each other on one radio channel. For data to be relayed from mesh node to mesh node, that node must repeat it in a store-and-forward manner. A node first receives the data and then retransmits it. These operations cannot occur simultaneously because, with only a single radio channel, simultaneous transmission and reception would interfere with each other (Figure 2).

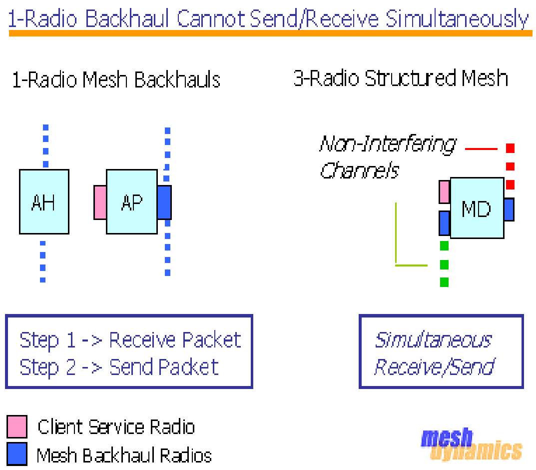

Figure 2: Single vs. Dual Channel Backhauls.

This inability - to simultaneously transmit and receive - is a serious disadvantage. If a node cannot send and receive at the same time, it loses ½ of its bandwidth as it attempts to relay packets up and down the backhaul path. A loss of ½ with each hop implies that after 4 hops, a user would be left with (½*½*½*½) = 1/16 of the bandwidth available at the Ethernet link. This is $1/(2^N)$ relationship defines the fraction of the bandwidth available to a user after N hops.

Third generation mesh products eliminate bandwidth degradation with a dual channel backhaul. There is been no measurable bandwidth degradation and this is the significant departure from both first and second-generation mesh architectures. Figure 3 shows live video from an IP camera part of a 9-hop mesh network at a ski resort. A single channel backhaul would be incapable of delivering this video feed beyond 1-2 hops: typical video bandwidth requirements would cripple the system. Latency and jitter would be unacceptable.

Figure 3: Live Feed from a multi-hop dual channel backhaul

Latency/Jitter Degradation on Single Channel Backhauls

Latency is inversely related to available bandwidth: thin pipes can provide only so much flow. Single channel backhauls suffering from bandwidth degradation also suffer from poor latency and jitter over multiple hops. This is primarily due to the need for a single radio to serve both backhaul and client traffic. The result is that most First Generation single-channel backhaul networks provide reliable video- or voice service over only one- or two hops. As a result many more costly wired or fiber Internet or intranet drops are needed to deliver adequate service, increasing the ongoing total cost of ownership.

Third Generation mesh products do not suffer from bandwidth degradation – they use multiple backhaul radios to obviate radio channel interference. Field tests indicate a latency of less than 1 millisecond per hop even under heavy traffic. Since latency is thus not a factor of traffic or user density, jitter (variation in delay) is also very low. This makes Third Generation products suitable for networks serving large number of users, demanding applications, video, and voice -- even simultaneously.

These capabilities of Third Generation wireless mesh networks make them especially useful where video surveillance is part of a metro/muni requirement, for expanding coverage into under-served areas with limited high speed wired or fiber infrastructure, and for border and perimeter networking, where bandwidth must be extended node-to-node as if in a long string of pearls.

Additionally, MeshDynamics multi-radio backhauls also incorporate VOIP concatenation for timely VOIP packet delivery. VOIP packets are small – typically less than 300 bytes but sent frequently, generally once every 20 milliseconds.

Networking protocols like CSMA/CA don't transport small and time sensitive packets well. The VOIP concatenation engine aggregates small VOIP packets into a larger packet for more efficient delivery. This aggregation takes place every 5-10 milliseconds. USAF tests (Figure 4) show overall latency is less than 10 ms + 1 ms per hop over a 4 hop network. Jitter is less than 1 ms. per hop.

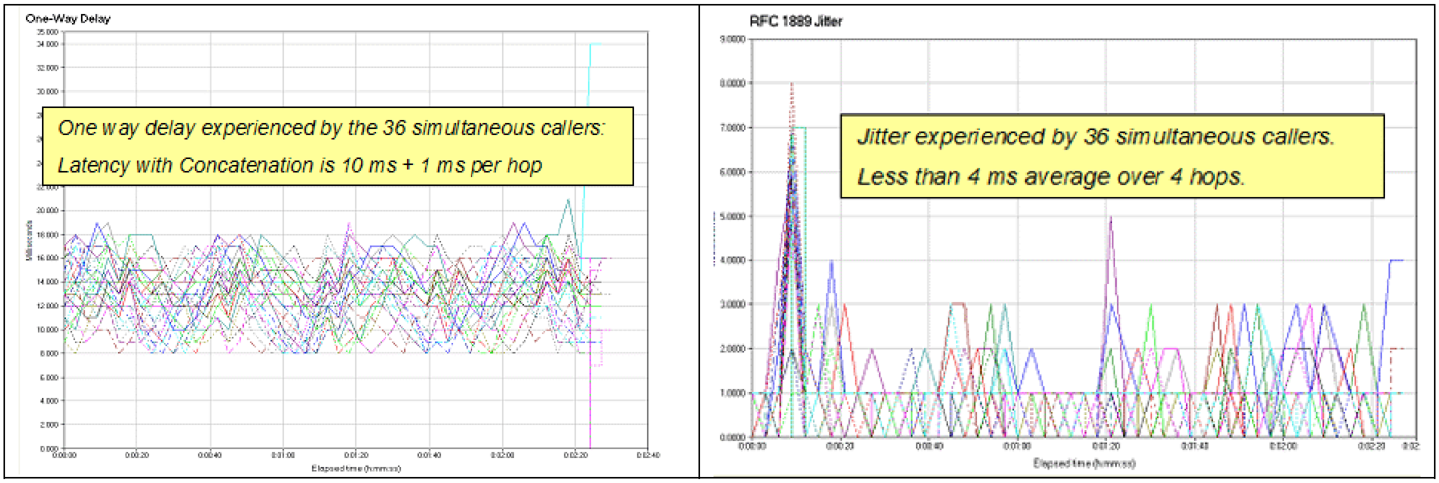

Figure 4a (left): Latency for 36 simultaneous VOIP calls over a 4 hop Mesh running VOIP Concatenation over the backhaul.
Figure 4b (right): Jitter for 36 simultaneous VOIP calls over a 4 hop Mesh running VOIP Concatenation over the backhaul.

USAF tests found comparable latency/jitter for single channel backhauls to be an order of magnitude higher. Some reasons for this:

- Bandwidth degrades with each hop. As an analogy to water pipes, the smaller the pipe, the slower the flow.
- With all radios on the same channel, there is compounded contention will packets fighting each other, all on the same channel.
- The overall efficiency of the CSMA/CA protocol used degrades exponentially as the number of clients on the same channel increase.

Frequency Agility through Distributed Intelligence

Wireless is a shared medium. Radios communicating on the same channel and within range of each other contend for available bandwidth. In single channel backhauls, there is one radio acting as the repeater between nodes: all backhaul radios must be on the same radio channel. The entire network is susceptible to channel interference/jamming. System performance is compromised. Figure 5 shows how the dual channel backhaul can switch channels to avoid debilitating external interference effects.

This first step of providing two-radio backhauls provides an obvious theoretical advantage over single-channel backhauls, but management of channel selection and interference avoidance becomes a challenge addressed in a number of different ways.

One of the first approaches applied to this problem was to segregate the backhaul links with hardware radio switching and sectored directional antennas. Since each sectored antenna "sees" only a narrow field of view, radio emissions from adjacent nodes do not create interference.

The limitations of this hardware-centric approach are the costs of sectored antennas and custom-developed radio hardware, the relative inflexibility of the system in dealing with perturbations and new external interference sources, and the complexity of site-surveying and installation. The alignment of the directional antennas must be precise and installation must include some manual determination of channel choices to maximize the efficient use (and re-use) of channels as well as a manual configuration of network topology.

A more-recent alternative approach utilized by MeshDynamics automates the channel-selection and topology-definition tasks by distributing dynamic radio intelligence in each node, in effect creating multiple "RF robots". Sophisticated algorithms allow each node to listen to its environment continuously to determine its relationship with neighboring nodes as well as extraneous and potentially interfering radio sources. Based on this analysis of the environment, an individual node selects the best channels to use to connect to the optimal nearby node for highest performance.

In this distributed dynamic radio intelligence approach, the network forms a tree-like structure emanating from one or more "root" nodes that have the wired or fiber connection to the Internet or intranet. As the branches of the "tree" radiate outward, eventually they become geographically distant enough from one another for nodes to begin re-using channels. This greatly increases the data-carrying capacity of the network, since it makes better use of the scarcest resource in an outdoor WiFi environment: the fixed number of unlicensed channels.

It may be obvious that this approach works with both sectored directional antennas and with omnidirectional antennas. The greatest flexibility comes with the use of omnidirectional antennas, since no pre-engineering of paths must be done. Each node ascertains the best connection and coordinates that connection with adjacent nodes through the periodic exchange of routing and other information.

Although the structure that results gives the appearance of a tree, the software intelligence in each node permits it to function exactly as the single-radio First Generation wireless mesh. A failure of any node prompts immediate coordinated reconnections around the network to bypass the failed node, as in the case of a traffic accident felling the light pole on which the node is mounted. When the missing node is returned to service, its neighbors recognize its presence and recalculate the best connections once again. This capability also makes additions and expansions to the network very straightforward, as new nodes may be simply configured with the proper security information, then powered-up. New nodes automatically are added to the network based on an exchange of information between the existing nodes, which are continually monitoring the environment.

The independent but coordinated "RF robots" in each node are also useful in detecting and avoiding external interference sources. Because 802.11 WiFi is an unlicensed medium, new independent and uncontrollable RF courses may appear in the network unpredictably and at any time. First-Generation wireless mesh networks are challenged by such sources since every device is sharing the same channel, but all may not be close enough to "hear" the offending point interference source. Hardware-oriented Third Generation networks may also be affected if an unexpected interference source appears in one of the sectored node-to-node paths.

Such an interference source is easily with managed distributed dynamic radio intelligence. Nodes close to the offending RF source may move in a coordinated fashion to a new and unused channel. By coordinating this movement, the impact to end users may be minimized. In extreme cases where a powerful interference source blankets an area, nodes cut off from communication with the rest of the network simply repeat the start-up process of listening for other nodes on non-interfered channels and the network rapidly reconverges to a stable state.

Quick and easy deployment without significant pre-engineering, automated and rapid avoidance of interference sources, and fast additions and reconfigurations of the network by the simple addition of nodes are all distinctive features of Third Generation architectures delivered through distributed dynamic radio intelligence as developed by MeshDynamics. An emerging application of this technology is in temporary and event-driven networking such as that needed for sporting events. An additional benefit of this software-based distributed intelligence is the ability to place a node in motion relative to other nodes in the network.

One use of mobility has been in security and other types of rapid-response applications. Nodes may be mounted on vehicles, even man-carried to new locations while remaining in communication with the overall wireless mesh network at all times. No management or other user interaction is needed, and with the use of omni-directional antennas, the mobile nodes may move in any direction around the perimeter or through the middle of the geographic area supported by the fixed wireless mesh node. Mobile nodes have even been mounted in unmanned aerial vehicles or tethered balloons to provide coverage of an area.

Another mobile application supported by MeshDynamics' distributed dynamic radio intelligence is support for networking in rail corridors. Mobile nodes installed in commuter trains link with a series of fixed nodes installed along the rail line. As the train moves, the mobile node is constantly listening to the environment. Typically, the signal from the fixed node being approached will be increasing while the signal from the fixed node recently passed will be decreasing. In a coordinated fashion, the "RF robots" in all three nodes coordinate the hand-off from one fixed node to another, assuring seamless connectivity for commuters in the railcar, whether the train is in motion or halted at a station.

Radio Agnostic Mesh

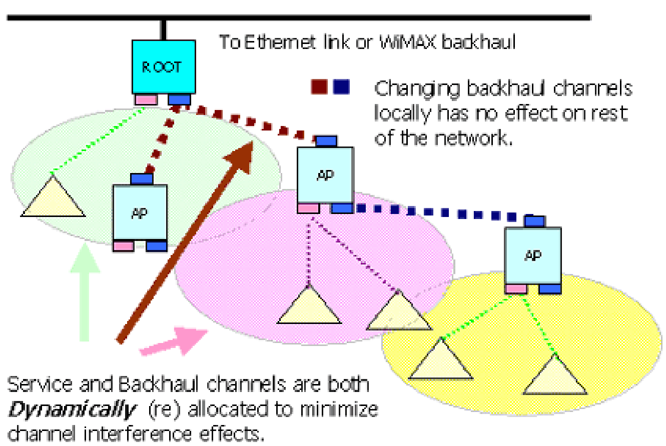

Figure 5: Channel Agility at every section of the mesh network.

Frequency agility is taken one step further in MeshDynamics Modular Mesh products. The mesh control "RF robot" software runs above the MAC layer of the radio: the same mesh control software supports radios operating on different frequency bands. Decoupling the logical channel-selection and topology-definition processes from the specific physical radio in this fashion delivers distributed dynamic radio intelligence benefits for current as well as emerging radio standards.

Figure 6 shows how this level of flexibility is supported. There are 4 mini-PCI slots on the board, two on the bottom and two on top. Each of the four slots can house a different frequency radio. This opens up some interesting possibilities including 2.4 GHz backhaul systems being part of a mesh with 5.8 GHz backhauls. Since the service and backhaul radios are distinct, it is possible to use a service radio to bridge over from a 5.8 GHz backhaul to 2.4 GHz backhaul – as shown in Fig 7. The 4325 Mobility Relay node on the bottom left has joined the mesh – even though the upper links are 5.8 GHz (blue) – through the service radio (pink).

Switching to another channel contains local interference at one section of the network. With one radio backhauls, this is not possible: the entire network is on the same channel and switching to another channel is simply not practical. The performance of single channel backhauls is therefore heavily compromised in RF polluted environments or under malicious attacks. Military field trials with dual channel backhaul have demonstrated that frequency agility ensures the mesh is running even with malicious RF interference- the backhaul radios (blue) simply switched to non-interfering channel.

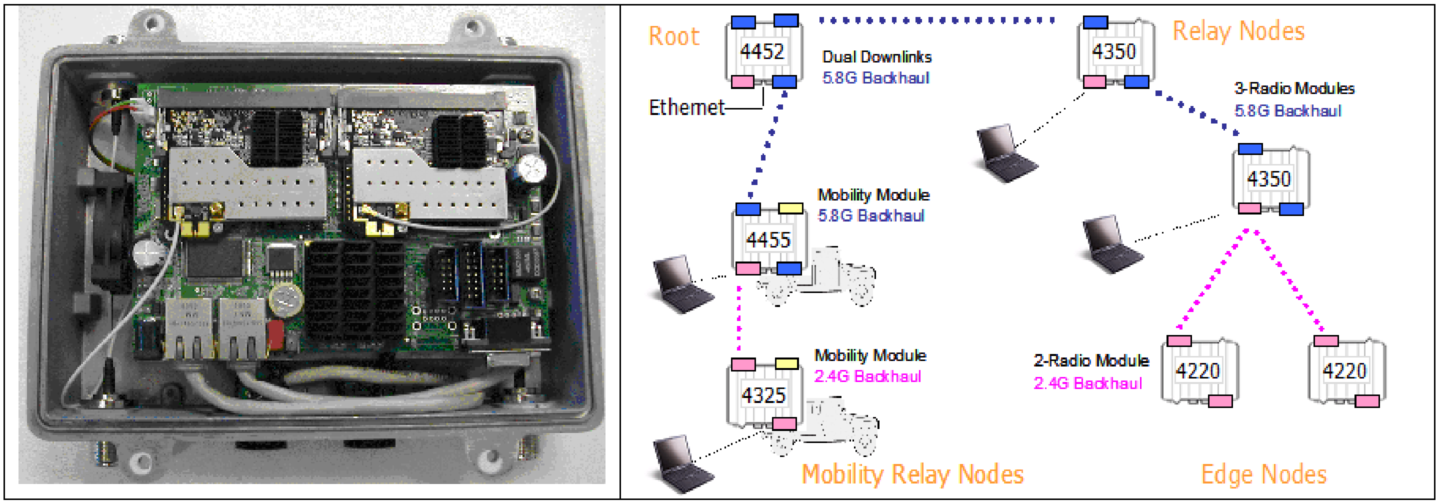

Figure 6, 7: MeshDynamics 4 radio slot module (left) supports 5.8 GHz and 2.4 GHz Backhaul Interoperability (right).

One advantage of this level of flexibility includes supporting longer range and lower bandwidth 2.4 GHz 802.11b radios with shorter range but more bandwidth capable 5.8 GHz 802.11a radios, all part of the same mesh network. The longer range enables the edge of the network – where bandwidth requirements are low – to be serviced adequately by 2.4 GHz edge/mobility nodes, (Figure 7). Bandwidth is thus distributed more efficiently and cost effectively managed: the node spacing is adjusted for the subscriber density based on the range of the radios used.

In the future, this flexibility will also permit the incorporation of new radio types such as WiMAX (802.16) as an adjunct to the 802.11 WiFi mesh. These potentially higher-speed links may offer a flexible way to inject bandwidth into the mesh along with a cost-effective distribution strategy for high-speed WiMAX links.

One additional benefit of decoupling distributed the dynamic radio intelligence software controlling channel selection and topology configuration from the radio hardware is the potential to incorporate commercial-off-the-shelf (COTS) radio and other modules. This substantially decreases time-to-market and greatly increases manufacturing scale for some components, reducing both development- and unit cost over the use of custom hardware development.

Greater Network Scale

Two elements allow MeshDynamics-based networks to scale to much larger overall sizes than networks based on competitors' products. The first is simply the nature of Third Generation technology. Because of the multiple radios in each node, it is possible to use separate, non-interfering channels for upstream links (toward the wired or fiber connection) and downstream links. The use of separate uplink and downlink creates the characteristic "tree" topology within a logical mesh network. Each of these uplinks and downlinks is a separate collision domain, vastly improving performance in terms of higher bandwidth, lower delay, and minimized jitter. The result is that much larger number of hops (node-to-node connections) may be supported, with production networks operating over twelve hops and twenty-hop networks being well within performance capabilities. This translates to high-performance networking over very large areas including hundreds or thousands of nodes, well beyond the capabilities of First- or Second-Generation technology. The difference between Third Generation and earlier technologies is somewhat akin to the vast difference between Ethernet switch and Ethernet hub technology: the segmentation of collision and contention domains creates more determinism in the network and permits much greater scale.

In addition, when using 802.11a for the backhaul protocol, the greater number of available channels permits re-use of channels as the network expands a distance beyond the wired or fiber connection. This re-use of channels is crucial to overcoming the scarcest limitation in the unlicensed wireless spectrum: the paucity of available channels. With virtually unlimited channel re-use in larger networks, they may scale to very large numbers of nodes without compromising performance.

The difficulty some users have experienced with earlier implementations of Third Generation technology is the requirement to manually define channel selection across the network. For very small networks this is less critical, but in larger-scale networks it rapidly becomes unwieldy. By contrast, MeshDynamics' distributed dynamic radio intelligence automates the process of channel selection and topology configuration. The distributed "radio robots" in each node automatically find other nodes in the network, choose the optimal channel and network uplink path, and continually monitor the environment to update these choices. This allows very large networks to be built without extensive pre-engineering, manual channel mapping, or antenna selection and aiming. Instead, new authorized nodes added to the network automatically join the existing network when power is applied, permitting rapid build-outs and expansion. Importantly, the network also dynamically avoids new interference sources and adjusts to other changes (such as a traffic accident downing a light pole supporting a mesh node).

In order to scale to larger sizes with adequate performance for high-demand data, video, and voice, both Third Generation and distributed dynamic radio intelligence are necessary. MeshDynamics provides the best choice for high performance without placing excessive pre-engineering and support demands upon network providers.

New Applications Enabled by Third Generation Wireless Mesh

Applications of earlier generations of wireless mesh technology were limited to a small number of node-to-node hops due to performance limitations, particularly the limited support for video- and voice-over-IP. This has created the perception that wireless mesh is suitable only for the delivery of basic networking such as casual web surfing. The emergence of Third Generation wireless mesh, with its inherent higher performance, is engendering many new and useful applications. Chief among these are applications requiring bandwidth to be extended across long distances by use of the wireless mesh backhauls alone. In these cases, it is impossible or impossibly costly to add additional wired or fiber network drops every two or three hops.

In one example, MeshDynamics wireless mesh nodes are being used along a national border to provide connectivity to distant locations in a long "string of pearls" configuration. In this case, there is no need to provide blanket WiFi coverage for client PCs or PDAs over the entire area, so the links between nodes are up to 14 miles in length. Mobile nodes mounted in security forces vehicles join the network dynamically and while in motion. Service radios in the vehicles provide connectivity for staff in the vehicles and operating nearby.

Similarly, a video surveillance application required extension of bandwidth into an undeveloped area with no installed high-speed infrastructure. Cameras are cabled directly to MeshDynamics nodes and linked back to a central site via many hops. Third Generation technology provides high performance and minimal delay and jitter to support the high fidelity video CODECs in use.

MeshDynamics' distributed dynamic radio intelligence technology, combined with Third Generation performance, is also key to a military application deploying sensors on moving combat vehicles with no fixed root wired or fiber connection. These vehicles are in constant motion in relation to one another, but information from sensors mounted on each must be brought together for threat analysis. The network topology must constantly and automatically adapt to the varying distances between vehicles while the mesh must deliver high performance with very low latency and jitter to permit the sensor data to arrive in a timely fashion. This combination of high performance and automated topology flexibility is enabling many other mobility applications that are not possible with other hardware-oriented Third Generation wireless mesh solutions.

Even in more traditional metro/muni applications, Third Generation technology is being used where earlier generations of wireless mesh have failed. In some localities, high speed internet infrastructure is not yet available from cable or telco providers or is prohibitively expensive. A single high-speed connection at the root node must be extended over a broad area using only the node-to-node connections for a backbone. This requirement for many hops resulted in the removal of earlier-generation wireless mesh and replacement with MeshDynamics equipment. The ease of deployment that comes with distributed dynamic radio intelligence can also be a benefit in these underserved areas where skilled RF engineers may be in short supply. And the reduced number of wired or fiber drops contributes to a lower total cost of ownership, permitting these networks to be deployed more quickly and more broadly.

Conclusions

New mesh requirements (more hops to cover larger areas, more efficient bandwidth distribution, better latency and jitter for Video and VOIP) have given rise to the Third Generation of mesh architectures. Third Generation multi-radio backhaul architectures deliver the higher bandwidth and more-deterministic performance necessary to meet these new requirements. While all Third Generation solutions deliver demonstrably higher performance, some custom hardware-oriented solutions come with requirements for extensive pre-engineering and offer limited capabilities for mobility and avoidance of interference. By contrast, software-oriented approaches based on distributed dynamic radio intelligence support frequency agility, automated channel selection, dynamic topology configuration, and radio agnostic meshes, providing more effective single framework solutions for larger scale and diverse application environments. This combination of features delivers higher performance for traditional metro/muni applications, but is also opening many new applications for wireless mesh.

About Meshdynamics

The MeshDynamics team began developing wireless mesh technology in 2002, with production unit shipments late in 2005. MeshDynamics' customers worldwide deploy our products in many applications: municipal (metro) networking; video surveillance, homeland security/defense; transportation; military; and public safety.

In addition to these sales of the standard products, MeshDynamics maintains substantial custom development activity for key defense applications. Much of the software pioneered for these demanding customers has been utilized as the basis for MeshDynamics' standard product line.

The MeshDynamics design is radio-manufacturer independent, allowing the rapid addition of new radio frequencies, radio system suppliers, and new technologies such as WiMAX and other emerging standards. MeshDynamics' technology is patented and patent-pending, developed by teams of engineers in the USA and India.