MIMO will clean up the mesh in your network

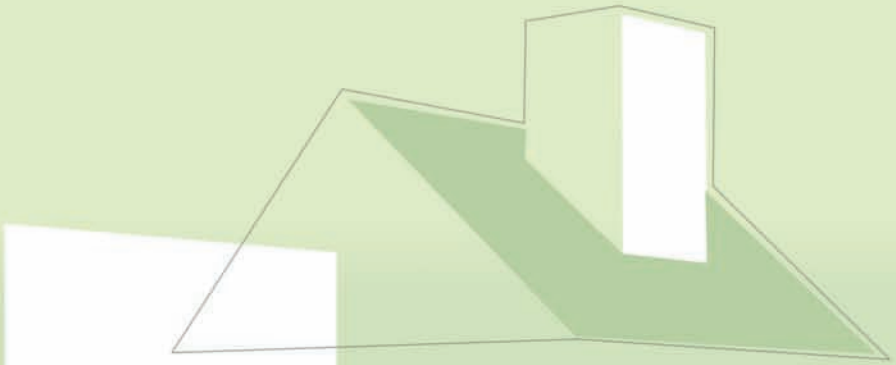

Mesh networks are evolving to support connectivity inside the home, with a number of CE companies looking at mesh WiFi as a way of improving the ability to provide ubiquitous home coverage. Philip Hunter explores the potential role that mesh can play in the connected home network

Wireless was too slow and unreliable for transmitting video within the home or anywhere else in the early days of WiFi, but the maturation of mesh and MIMO (Multiple Input Multiple Output) technologies is changing all that. By the time home networks are widely deployed, WiFi is likely to be the most popular option for buildings that do not already have a wired network installed. The latest HomePlug AV standard, enabling transmission of data and video over the mains wiring, also avoids the need for additional cables, but only supports portability between sockets, and not full mobility as required by the growing number of battery powered or rechargeable devices.

Mesh comes in various guises. It first evolved in trunk telephony networks to allow re-routing in the event of congestion, or failure of either nodes or links. In a full mesh, every device is connected to every other, so in a wired network, the amount of cable required increases geometrically with the number of nodes - a three node network requires just three links, while a six node network requires 15 links. For this reason, wired networks with large numbers of nodes tend to be just partially meshed, with each device connected to a few others, sometimes only the immediate neighbours.

Within the home, meshing could be deployed in principle across a wired network but there would be little point since it would not solve any bandwidth problems, which are not really the issue here, while the small risk of node or link failure can be tolerated. In any case, wired home networks will often utilise existing physical networks and therefore the physical topology is already set, as with power or telephone cabling. The story of mesh in the home is, therefore, almost entirely about wireless.

Extending the mesh with MIMO

Wireless mesh technologies, however, evolved first for metro or local loop networking and are only now being considered for deployment in the home. Meanwhile, MIMO has given home networks a boost by exploiting the varying paths radio signals tend to take between points when there is no direct line of sight. MIMO tackles two problems. Firstly, it reduces the incidence of communication blackspots where signals previously did not reach, and secondly, it improves Quality of Service (QoS) for real time applications such as video, by providing an alternative path in the event of transmission interruptions.

"MIMO leverages what has traditionally been a liability for radio, which is the multiple paths with signals bouncing around," says Frank Hanzlik, managing director of the WiFi Alliance, which formed a group in late 2006 looking at mesh networking capability and applications in the home. With appropriate signal processing as in MIMO networks, this multiple path effect can be exploited to provide redundancy and network penetration.

But MIMO cannot guarantee to reach all corners of every home on its own. In some cases, the physical barriers to communication between any two points are too great for signals to traverse via any path, as Dave Park, VP of product marketing at BelAir, a provider of mesh networks for hotels, campuses and metropolitan areas, points out. "The main challenge in the home is the propagation environment and mesh could help, for example, reach into the basement," he says. In this case, a device - say located at the top of the house - might not be able to transmit signals directly to one in the

basement, but could do so perhaps via an intermediate device in the middle of the house. MIMO, in turn, would ensure that reliable links with adequate QoS are maintained between devices within the mesh.

Fluidity is key

With a wireless mesh, there are no wires and therefore no cost associated with the physical layer, unless this uses licensed spectrum, which would not be the case in the home. However, the allocation and use of spectrum has been a major issue in the evolution of wireless mesh for metropolitan networks and as an alternative to DSL or cable as a local loop for delivering broadband services. In such external networks, the role of meshing lies largely in providing coverage rather than QoS or reliability, for there will often be line of sight between nodes and the objective is to replace wires for interconnecting access points. "Meshing provides integration between wireless access and wireless transport," notes Park.

Early wireless networks, including WiFi hotspots, always comprised access points that were linked back into delivery infrastructure by wired connections. Client devices such as laptops or PDAs would then connect over wireless to the closest access point. The aim of mesh in the first instance is to makes these wired transport connections wireless as well. At present, the wireless transport is confined to the local or metropolitan area, but in principle it could be extended over larger distances, although currently it is questionable whether it would scale well or be cost effective for transmission over remote areas where there would be relatively little access. Within these networks outside the home, mesh has initially been implemented only between fixed access points, providing coverage between roaming devices inside the zone.

To have a mesh, there has to be a minimum of three such points, whereas in the home there is usually just one access point serving all client devices. Therefore, a home mesh would be implemented in the WiFi clients. This changes the game, for while some of the clients might be relatively tied to one location, for example a desk bound PC or TV, others, such as laptops or even mobile handsets, would be roving. Furthermore, devices will not always be switched on or even present so the home mesh has to be totally dynamic and fluid, able to reconfigure itself on the fly to accommodate the shifting and disappearing nodes.

Support for such fluidity has not been necessary for metro WiFi applications, typically comprising access points installed on poles or rooftops, but some vendors have adopted mobile-mesh anyway to provide the added flexibility and be ready for the home market. One such vendor is PacketHop. "Unlike fixed-mesh networks, ours is highly mobile and the network adapts as users come and go within range of the network," says Kevin Payne, PacketHop's director of marketing communications. "And because it is a mesh network, it is also highly resilient. If one node leaves or is congested, traffic is routed elsewhere throughout the mesh."

Such a mobile mesh poses particular challenges for video, because the reconfiguration or 'self-forming' process as devices move around inevitably takes time and imposes latency constraints. According to Payne, PacketHop has addressed these issues by ensuring that all devices within range of each other maintain connection so that when one drops out, traffic can be quickly rerouted. The network maintains the optimum number of video streams at all times.

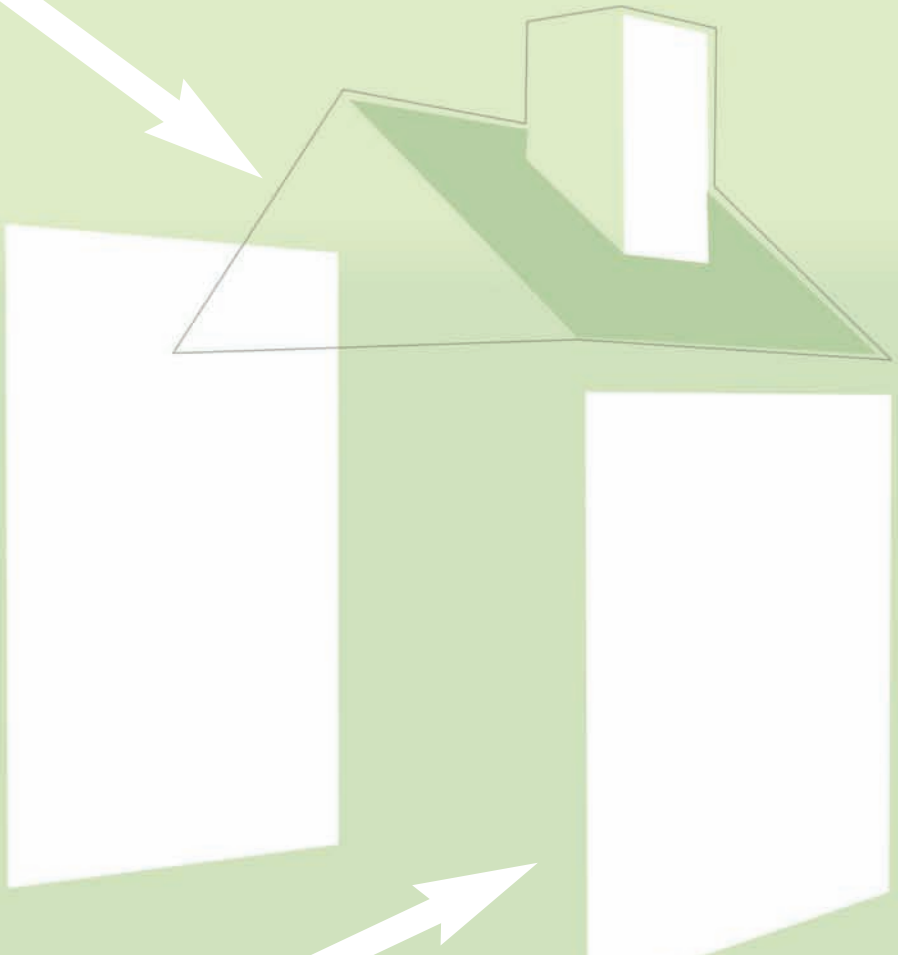

"Devices will not always be switched on or even present, so the home mesh has to be totally dynamic and fluid, able to reconfigure itself on the fly to accommodate the shifting and disappearing nodes."

"The IEEE 802.11s amendment is trying to get the 802.11 mesh network working at layer two and this seems a very hard issue to cope with, because it involves extensions that are very difficult to implement at hardware level."

The support for video has been driven by PacketHop's customers in security and surveillance at the metro level but will be just as relevant in the home, says Payne. Indeed, mesh will only come into its own within the home for high bandwidth applications such as video, Payne asserts. "For example, if you were trying to stream video between computers, meshing would offer some advantages because the mesh creates more bandwidth as more nodes/ users are added," he says. "So potentially as more devices, such as home entertainment systems, become WiFi enabled, meshing may offer some advantages."

Meshing could also reduce the chance of temporary breaks in service when a device roams around. A mobile handset might be moved into a spot where it is out of range of, say, a home PVR streaming video to it, even with the help of MIMO. But it might be closer to another device, which in turn is still within range of the PVR. So while mesh extends distance at the metro level, within the home it could have the effect of joining up lines of communication.

Like MIMO, mesh has the effect of turning conflict between signals into cooperation, as is noted by Roberto Arcadu, general manager of WiNext, which has recently introduced a home version of its mesh solution based on its New Wireless Auto-configurable Device (NAAW) broadband access system, which enables any WiFi device to participate in a mesh. "A mesh network is more stable as all the nodes cooperate instead of conflicting, and the result is a better quality signal and a lower consumption of energy and less radio pollution," says Arcadu.

Moreover, some of the problems facing mesh networks over a larger geographical scale do not arise within most homes. "Because the range is limited in the home, performance may actually improve," says PacketHop's Payne. "A mesh network in the home wouldn't be likely to face the interference or usage challenges that a metro WiFi would face."

Which layer for in-home mesh?

Routing is also simpler within a home mesh, which has an impact on the debate over whether it is best to do at layer two or layer three. This debate has resurfaced at different levels for almost two decades, ever since local Ethernet networks started to be interconnected into more complex wide area mesh-like configurations. Layer two simply refers to Ethernet itself, and the process of transmitting data between nodes on a fixed point-to-point basis with only limited support for re-routing in the event of failures or congestion. Layer three evolved to support more complex routing within mesh networks, and therefore comes with the fundamental ingredients for wireless mesh networking already built-in. For this reason, some wireless mesh vendors have adopted the layer three approach, arguing that layer two needs substantial extension to support mesh networking.

But there is also a competitive dimension, with layer two standards being promoted by the IEEE and layer three by the IETF. "The IEEE 802.11s amendment definition is trying to get the 802.11 mesh network working at layer two and this seems a very hard issue to cope with because it involves extensions that are very difficult to implement at hardware level," says WiNext's Arcadu. "On the other hand, IETF already proposes a complete mesh solution laying on OLSR (Optimized Link State Routing) or AODV (Ad hoc On Demand distance Vector) routing protocols. This is a de-facto standard that many mesh network providers around the world are already adopting to deploy systems in large WiFi covered areas."

Yet within the home, some vendors regard use of such protocols as unnecessary overkill, given that there is not that much real routing to be done. "Layer 3 in the home would be a big mess," insists Matt Holdrege, technical director of the wireless mesh pioneer Strix Systems.

Integrating home WiFi networks

But this raises another issue, which is how home WiFi networks will integrate with the local loop or metro network. If the latter are also wireless mesh based, then it might after all make sense to support layer three routing for end-to-end transport. A likely compromise, according to BelAir's Park, will involve layer two within the home, but then allocating IP addresses to cross the external network at layer three.

There is the option, however, of supporting the home mesh as an extension of the public or external network, as has already been tried with some community projects. The user would still appear to have a home network but in practice this would just be part of a much larger mesh. For this to work, mesh networks need to scale to very large numbers of nodes without degradation in QoS. On this front, Microsoft is among the pioneers, having developed a new routing protocol called Virtual Ring Routing (VRR) at its research laboratory in Cambridge, UK.

This is an interesting development because it takes account of the limited range of each hop in a wireless mesh, with the whole network inevitably comprising a large number of hops if it has any geographical scale. "This takes a new approach to routing by merging ideas for peer-to-peer overlay networks and traditional routing protocols," says Peter Key, head of the system and network group at Microsoft Research (MSR) in Cambridge. "Meshes have already been shown to scale to the order of 100s of nodes, and we have one of the world's largest test-beds in our lab at Cambridge consisting of over 110 machines."

However, this is still a long way from the size of a full blown metropolitan network, emphasising that these are early days. As the WiFi Alliance's Hanzlik agrees, it is too soon to tell whether the community approach will prevail, or whether consumers will prefer to have their own private home mesh. "There will probably be different go-to-market models," he says.

The whole issue of spectrum allocation then arises, with the intriguing possibility of convergence or integration between the licensed spectrum used in public networks, and unlicensed spectrum as used by WiFi at the home or neighbourhood level.

Amid all these debates, there is agreement at least that the home mesh will come to market one way or another. **CSI**