

1 Mesh overview and terminology

The Internet is now firmly part of our everyday life. We perform many common tasks on-line, such as banking, grocery and gift shopping and the purchasing of travel or cinema tickets. Plus we get a growing portion of our entertainment from on-line sources: entertainment and social networking are two of the largest growth areas. We have seen the beginning of basic quality video from, for example, YouTube and the development of social networking sites such as MySpace and FaceBook, which have been enormously popular, especially amongst younger generations of consumers. If we are to continue in this trend of doing more on-line, our need for bandwidth will increase. And in future we might expect to generate appreciable content ourselves, for upload onto the Internet, as well as to continue to download content. But that is not all; our need for Internet availability and quality will also increase.

It would be very convenient if such future Internet access were also wireless, with the near ubiquitous service we are used to from cellular phones. However, building a new network to achieve this, or upgrading an existing network to support this, would mean installing or changing a great deal of infrastructure. What then if a method existed which promised improved Internet access with fewer requirements for new infrastructure? This is widely advertised as the domain of the mesh network.

This chapter begins with a top-level introduction to mesh networking, then looks at how meshes may fit into the larger telecommunications infrastructure, before moving on to classify and explain the basic properties of a mesh. Finally the chapter closes by bringing together the key issues for mesh networking and by linking these issues to the structure of the book. This includes four hypotheses of how meshes might be expected to bring networking benefits. The hypotheses are real, having been taken from a growing mass of mesh-centric literature at the time this book was written. Testing these hypotheses will form a useful basis for investigation and they will be revisited often as we progress through the book.

2 MESH OVERVIEW AND TERMINOLOGY

1.1 What is a mesh?

It is perhaps easiest to begin examining mesh networks by first taking a small step backwards, by reviewing how cellular and wireless local area networks (LANs) work, before highlighting the similarities and differences embodied in the mesh approach.

A cellular network, as its name suggests, consists of many cells of radio coverage, each having a base station near its centre which transmits radio signals over a wide area, for example several kilometres in diameter. The user's device is a small handheld unit of lesser complexity and capability than the base station. Where the coverage of one cell falls off, it is picked up by adjacent cells. In this manner, a large area is covered by a honeycomb of cells. Clearly the advantage is the promise of continuous coverage. This is a major advantage, but there are some downsides: a new network must be planned, plus, ideally, all cell sites should be rolled out simultaneously over the whole coverage area, but this means a large upfront cost for the network operator.

Despite best efforts, some black spots may occur where the user may be shielded from the radio signal by an obstruction. Whilst one way around this would be to install an additional small cell, too many of these would increase the cost of the network infrastructure and hence adversely affect the operator's business model. In reality a compromise is drawn whereby a typical network availability approaches 100%, but does not guarantee to please absolutely all the people, absolutely all the time.

The overarching method used to ensure good coverage is to choose a carrier frequency which propagates well across many terrain types and has good building penetration characteristics. This will jointly maximise coverage and capacity and minimise the number of cells.

It is no coincidence therefore that cellular systems worldwide centre around a very similar band of frequencies. Such spectrum is highly valued and already allocated and assigned in most countries. It is worth noting that cellular systems were designed primarily for voice communications, which are bursty in nature. The amount of spectrum they have is thus not ideally suited to modern multimedia communications, which may involve the streaming of delay sensitive data over a relatively long period, for

example to watch a video. The future evolution of cellular systems is therefore aimed towards better supporting multimedia applications.

Let us next look at the situation for wireless LANs. A wireless LAN, as its name suggests, is mostly about local area coverage. The objective is usually to cover a workplace or a home. Commonly, wireless LANs have a design objective to cover up to a 100 metre radius around a wireless access point on a wired LAN. The access point and base station concept are similar, but the distances under consideration are very different. One immediate benefit is that the type of spectrum required is different, specifically that the propagation characteristics need not be so good for wireless LANs. Higher frequencies are adequate to cover shorter distances, leading to typical wireless LANs operating at higher frequencies than typical cellular systems. Here the spectrum is in less demand. In fact many wireless LANs operate in spectrum which is licence exempt, meaning system costs are lower. But this is a double-edged sword; unlicensed wireless LANs do not have exclusive use of their spectrum as do cellular systems. This leads to the need to design the wireless LAN to be tolerant of interference, which has costs in terms of equipment and system operational efficiency. In the extreme, efficiency may drop to zero if congestion occurs, as it may do in any unlicensed system.

There are at least two more reasons why a wireless LAN may be simpler to implement than a cellular system. In general, there is no concept of cell to cell hand-over in a wireless LAN, nor does the wireless LAN offer much more than a raw data link between users; it is up to additional protocols such as IP and TCP to establish routes and transport the data reliably. It is worth noting that wireless LANs were designed primarily to transmit data rather than voice. However, a design constraint of transporting bursty data was assumed, which was realistic at the time. This means that whilst wireless LANs have good capacity for large amounts of data, they typically do not cope well with modern multimedia communications. In other words, whilst wireless LANs were designed primarily for cases where demand is higher at peak times, such peak times were not expected to be of very long duration. The future evolution of wireless LAN systems is now directed towards better support for multimedia applications, much like the evolution roadmap for cellular systems.

4 MESH OVERVIEW AND TERMINOLOGY




Figure 1.1 A spider's web, an example of a mesh.

Turning, at last, to the mesh network, once again we can infer its prime working attribute from its name. Meshes are all around us in one form or another – think of a spider's web, Figure 1.1, or the grid pattern of streets in a downtown area (and there is an application clue...). Imagine that at each material intersection there is a node. What these examples have in common, for our purposes, is two-fold:

1. there is no main node;
2. it is possible to reach any other node by traversing a number of intermediate nodes.

Immediately we can see that a mesh architecture is quite different from a cellular or wireless LAN architecture. All nodes are equal so there is no centralised control and, therefore, each node must participate in networking as well as be a source or sink of traffic. Rather than a single hop to a base, multi-hopping amongst nodes must be a common capability.

In fact this all brings the promise of great flexibility, particularly when we wish to create a new network, or expand an existing one. By way of example, consider that we wish to network five people in their house and garden, an increasingly common task. Let us try to do this firstly following the access point or cellular principle and secondly by the mesh principle. We will keep the example extremely simple.

Let us say that the range of all wireless units just exceeds the maximum dimension of the house. Then the situation for cell-based and mesh is equivalent; each approach can cover the house. Let us now say one person moves to the far end of the garden which is very long, much larger than the house dimensions. The cell based scheme cannot cope with this; the radios are simply out of range. The mesh based system can cope with one proviso, that a third person moves in between the house and the end of the garden. All units may now communicate as before. The unit at the end of the garden simply multi-hops, using the third person's unit as a relay. It is thus easy to see that multi-hopping can cope with distance, and we should bear in mind that this must mean it can also cope with clutter by hopping around obstructions.

Finally, let us say all the users and their wireless units move to a new house. Unless they take the base station with them, the cellular scheme clearly will not work. On the other hand, the mesh scheme works as before. This is the quality of meshes to work without infrastructure.

Whilst this example may be simplistic, it is also realistic. The principles remain the same when we look at more involved examples of where and how meshes might fit in with existing telecommunications systems, in the future.

1.2 The role of mesh in future networks

In order to understand the wider role of mesh networks it is necessary to place them in the context of an overall communications environment. Like many in this field, in the future we believe this is likely to consist of a wide range of different wireless communications systems, connected to a single core network based around IP packet switching.

We therefore see this future as an evolutionary integration of the cellular approach with the WLAN approach as depicted in Figure 1.2, where proprietary interfaces and protocols have largely been removed.

In the scientific literature, this is sometimes referred to as 'B3G', meaning 'beyond 3G'. Such terms are usually proffered by the existing cellular focus groups, but WLAN parties also have a similar vision – perhaps this should be called 'beyond WiFi'. Thus B3G and BWiFi are the same integrated vision and as such must surely delay or even remove

6 MESH OVERVIEW AND TERMINOLOGY




Figure 1.2 Future mobile integrated vision using an IP core.

the need for a wholly new and bespoke fourth generation (4G) network. In other words, we think 4G will look like Figure 1.2. Meshes are expected to provide a complementary access route into the core, alongside WLAN, 3G etc. In other words, mesh will be an additional access technology rather than a replacement. Whilst the foregoing may seem self-evident, there have been some who have thought mesh would be more of a unilateral revolution than a component of evolution.

It is worth looking at some of the typical types of networks presently used to link into the core of the network. The core is used to connect the users to the Internet via the Internet service providers (ISPs). Content providers may be particularly closely aligned with ISPs in some future business models. The core supports access at very many different data rates from different devices, which is a good match for the granularity of service offered by TCP/IP. Example technologies include the following.

- PAN – personal area network. This may for example use Bluetooth and may even be mesh-like in its structure. The speed is currently under 1 Mbps (although 100Mbps is predicted) and the range is short, e.g. 10 metres. Access points are needed.
- GSM/3G – second and third generation cellular mobile. Data rates will initially be under 2Mbps with the average around several hundred kbps.

The range is high, on the scale of several kilometres. It relies on a deployed infrastructure of base stations.

- DSL – digital subscriber line. This will typically be a point to point connection running over fixed copper pairs carrying Ethernet or ATM frames. The speed is currently around 8Mbps, depending on distance and network loading (contention).
- WLAN – wireless local area network. This provides potentially the fastest access towards the core; up to 54Mbps is available, with 11Mbps being widespread today. Its range is around 100 metres. Some WLANs enable a choice of whether to use infrastructure or peer-to-peer, but the majority are using an access point infrastructure.
- Mesh – presently only very sparsely deployed in the ‘early adopter’ market. Its performance capability is to be investigated within this book. It is potentially high speed, with good coverage and with no need of infrastructure.

One key point from Figure 1.2 is that there are two ways at present to obtain relatively high speed wireless access to the core – via cellular and via WLAN. They have different performances in terms of quality, speed and range, but also critically in terms of cost. The cost benefits to the user from having a dual mode handset for attachment via a WLAN at close range and via cellular at longer range can be very attractive. However, such ‘fixed-mobile convergence’ presents a real challenge to operators who may well have to change their business model and adapt their technology. But mesh is relatively new to the vision and may enable radically fresh approaches to the challenge of high speed wireless access for the future.

1.3 How do meshes work?

Before we investigate mesh in detail we need to understand more closely what a mesh network actually is. This section discusses and defines the nomenclature and methodology used for mesh networking. It continues by clarifying the meanings of the terms mesh and ad hoc.

8 MESH OVERVIEW AND TERMINOLOGY

1.3.1 Forms of mesh

There are broadly three basic mesh type architectures, to which we shall refer throughout the discussions within this book. The three types are introduced here and described in detail in Section 3.2.2.

- Pure mesh – all traffic in a pure mesh is intra-mesh, i.e. the mesh is isolated. All traffic is relayed via a single node type, i.e. the user node.
- Hybrid mesh – a pure mesh with a hierarchy of node types, in order to improve efficiency via the introduction of backbone routes. In other words there is a further network overlaying the mesh, which exclusively takes longer distance traffic. In wireless meshes, such a hierarchy of routing layers may be created simply by using additional dedicated radio channels, or bands.
- Access mesh – a mesh with a hierarchy of node types, as above, but additionally where considerable traffic is extra-mesh. In other words, the overlay routing network also has gateways to other, external networks such as the Internet.

We shall see later that traffic flow and hence the most appropriate mesh architecture depends on whether the content to be accessed resides inside or outside the mesh. In other words, the type of mesh required in a given situation is driven by the needs of the user and the application. It is worth noting from the outset that most of the early published research was funded by the military to look at pure meshes and does not always translate well to public telecommunications requirements, which generally require an access mesh.

1.3.2 Planned versus ad hoc

A second distinction is concerned with the design rationale of the network. A planned network such as cellular has a predetermined maximum level of users and protected cells in which they may operate. The benefit of this is that interference is limited and it flows from this that guarantees of the quality of service delivery can be made. The downside is that infrastructure is needed; in other words the operator must first make a provision for everywhere the user requires a service.

An unplanned network allows ad hoc connections. ‘Ad hoc’ literally means ‘for this purpose’, i.e. a temporary arrangement as and when required. The benefit is that no infrastructure is needed and the users themselves may extend the area of coverage – although we shall see later that must involve a performance trade-off, if all else remains equal. The downside is that without planning, there is no control over interference effects from other users. Hence an application’s performance is beyond the control of any one party, therefore absolute guarantees over the quality of the delivered service cannot be made.

1.3.3 Characteristics of an ad hoc pure mesh network

Table 1.1 briefly summarises the characteristics of a pure, ad hoc mesh network. These are the characteristics which will be examined in later chapters.

1.3.4 Characteristics of an access mesh

Moving up in complexity, a hybrid mesh is similar to the type of mesh described in Table 1.1, except that some nodes will additionally have independent network connections, enabling them to connect to an internal routing backbone. One logical step beyond this would be to add the ability for some nodes to be gateways to external networks. Of course this turns out to be exactly the configuration which is needed to provide Internet access to a group of users. For this reason, we refer to this mesh type as an access mesh. The access mesh is the most important mesh type for the rest of this book.

1.3.5 Meshing versus multi-hopping

We have already seen a simple example of home wireless communications, where multi-hopping was useful to increase coverage by increasing distance or hopping around obstacles. This begs the question of the difference between meshing and multi-hopping.

10 MESH OVERVIEW AND TERMINOLOGY

Table 1.1 *Characteristics of ad hoc mesh networks*

Ad hoc	Unplanned. Therefore the coverage and interference environment is uncontrolled, which is the exact opposite of the typical cellular case. This directly raises quality of service issues.
No separate infrastructure	All functionality is provided and carried out within the mesh. This includes any power control, routing, billing, management, security etc. There is no centralised equivalent to the base station or security and authentication centre of, for example, cellular networks. (Not true for an access mesh, see below.)
Mobility	Nodes are free to move and even disappear. The network interconnections may thus be very dynamic.
Wireless	In order to support mobility or avoid infrastructure, wireless operation is required. This could mean radio or optical, but this book concentrates on radio. Radio links are lower quality than wired links, packet loss in radio is ‘normal’, whereas on wired connections, loss is equated to congestion. Transport protocols (as developed for wired networks) may thus have the ‘wrong’ reaction when used on radio networks.
Relay	All nodes may be required to relay information for other nodes. This will lessen the bandwidth available to each node user.
Routing	All nodes will be required to participate in a routing protocol. This may be either proactively by maintaining up to date tables or reactively by creating routes on demand (which may also be cached in tables). Routing creates an overhead, which will depend on the protocol, the traffic and the mobility of the nodes.
Multi-hop	A corollary of relay and routing, multi-hop is an enabler of coverage, especially in the cluttered environment.
Inhomogeneity	Not all nodes need be equal, beyond the subset of capabilities needed for basic mesh operation. Some nodes may have additional network connections (external connections in the access mesh case, see below).