Essentials of Wireless Mesh Networking

CAMBRIDGE WIRELESS ESSENTIALS SERIES

STEVE METHLEY

CAMBRIDGE www.cambridge.org/9780521876803

This page intentionally left blank

Essentials of Wireless Mesh Networking

Are you involved in implementing wireless mesh networks? As mesh networks move towards large-scale deployment, this highly practical book provides the information and insights you need. The technology is described, potential pitfalls in implementation are identified, clear hints and tips for success are provided, and real-world implementation examples are evaluated. Moreover, an introduction to wireless sensor networks (WSNs) is included. This is an invaluable resource for electrical and communications engineers, software engineers, technology and information strategists in equipment, content and service providers, and spectrum regulators. It is also a useful guide for graduate students in wireless communications and telecommunications.

STEVE METHLEY has over 20 years' experience in telecommunications and data communications innovation, having led teams in the laboratories of British Telecom, Hewlett–Packard and Toshiba. Currently a consultant based in Cambridge, UK, his work includes technology, regulation, business strategy, socio-economic analyses and futurology for a range of clients from start-ups to large global corporations.

The Cambridge Wireless Essentials Series

Series Editors
WILLIAM WEBB *Ofcom, UK*
SUDHIR DIXIT

A series of concise, practical guides for wireless industry professionals.

Martin Cave, Chris Doyle and William Webb *Essentials of Modern Spectrum Management*
Christopher Haslett *Essentials of Radio Wave Propagation*
Stephen Wood and Roberto Aiello *Essentials of UWB*
Christopher Cox *Essentials of UMTS*
Linda Doyle *Essentials of Cognitive Radio*
Steve Methley *Essentials of Wireless Mesh Networking*

Forthcoming
Albert Guiléni Fàbrigas *Essentials of Error Correction for Wireless Communications*

For further information on any of these titles, the series itself and ordering information, see www.cambridge.org/wirelessessentials.

Essentials of Wireless Mesh Networking

Steve Methley

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore,
São Paulo, Delhi, Dubai, Tokyo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521876803

First published in print format 2009

ISBN-13 978-0-511-57729-1 eBook (Adobe Reader)

ISBN-13 978-0-521-87680-3 Hardback

Contents

Preface

Wireless mesh networking is a hot and growing topic, still in its infancy in some ways, whilst already shown to be capable in others. From a military beginning, mesh networks moved to civilian use and are now being deployed worldwide as both local area networks (LANs) and metropolitan area networks (MANs). However, these deployments are still 'leading edge' and it is not yet clear what the most enduring applications of mesh will be – particularly as the market moves from early adopters towards widespread take up.

Some of the claims for what a mesh network may deliver have been very ambitious to say the least. In this book we investigate such claims versus the real qualities of mesh networks and identify the key time scales and drivers for the challenges involved with making meshes. Throughout the book we attempt to keep mathematics to a minimum. Where an equation is shown, it remains practical to follow the flow of the book without needing to understand the maths fully.

The book takes a very pragmatic but balanced approach to the issues. We are particularly interested in meshes with an external access capability, for example to the Internet. We supply a technical assessment of mesh and multi-hop networking, highlight the attractions, identify the pitfalls, provide clear and concise hints and tips for success – summarised inside the back cover – and finally evaluate some real-world examples of good mesh applications. These include wireless cities, community networking and vehicular ad hoc networks (VANETs). Wireless sensor networks (WSNs) are another important application of mesh techniques with their own unique challenges, and these receive their own chapter.

We conclude that, although some of the claims for what a mesh may deliver have been exaggerated, the real qualities of a mesh network when directed to a suitable application can still make mesh the best approach.

Acknowledgements

The author is very pleased to be able to thank past and present colleagues for many vigorous and entertaining technical discussions. Two former colleagues, Malcolm Crisp and James Newman, deserve special thanks for the many animated brainstorms we had together whilst deliberating the issues in Chapter 4. I was fortunate enough to be able to talk to many experts including Saleem Bhatti who expanded my appreciation of networking issues, especially for sensors. Other experts included Peter Massam who made many useful comments on modulation schemes, Peter Ramsdale who is now sadly missed, Frank Rowsell, Chris Davis, Nee Joo Teh, Stuart Walker, Ahmad Atefi and last but certainly not least William Webb, our book series editor, who provided many very helpful suggestions for the manuscript. Finally, I also belatedly want to thank Peter Cochrane who gave me so much encouragement in my first role all those years ago at BTRL.

Thanks are also due to my wife for her support at all times without exception and to our two sons for having such boundless energy which is a constant source of inspiration.

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.

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.

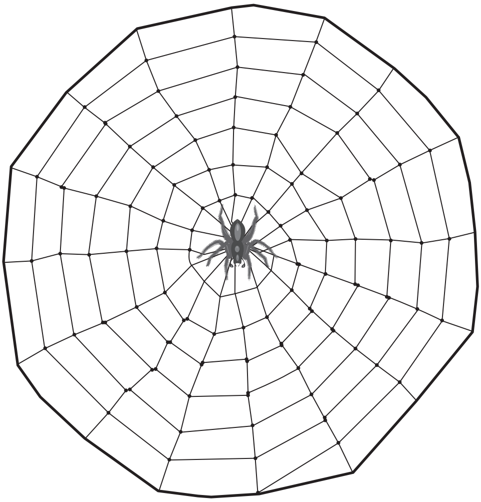

Figure 1.1 A spider's web, an example of a mesh.

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

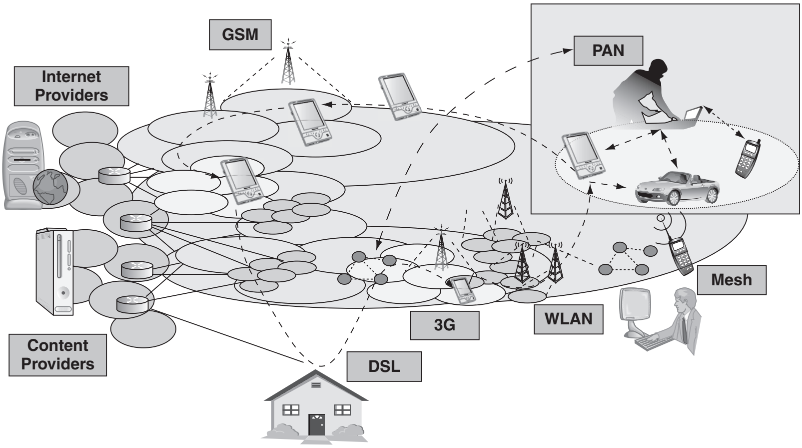

Figure 1.2 Future mobile integrated vision using an IP core.

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 100 Mbps 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 2 Mbps 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 8 Mbps, depending on distance and network loading (contention).
- WLAN – wireless local area network. This provides potentially the fastest access towards the core; up to 54 Mbps is available, with 11 Mbps 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.

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.

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).

In fact the terms are not universally well defined and mesh is often used generically. Where it matters, in this book we will treat the difference as follows. Meshing may simply be thought of as multi-hopping with active route diversity. To use metaphors, a multi-hop network may be thought of as nodes within a tree and branch structure, whereas a mesh is more like nodes within a spider's web. In other words, in a mesh a given traffic flow may be split over two or more routes to the destination, whereas in a multi-hop network there is a single routing at any point in time. Nonetheless the multi-hop network preserves the ability to route around obstacles. This common, important property of hopping around obstructions will surface many times within this book and is perhaps the most enduring benefit offered by both mesh and multi-hop networks.

Looking a little deeper, it may be realised that the difference between meshing and multi-hopping can become quite blurred. If a multi-hop network may reconfigure another route to the destination, is that not a mesh in effect? One remaining difference which may matter becomes the fact that meshes may have two or more routes active simultaneously – but that brings problems along with benefits (problems may occur at the TCP level, which, in practical implementations, does not suffer reordered packets gladly).

However, a full mesh has redundant routing out to the network edge – multi-hopping tree and branch does not do this, neither does the mesh-star approach of ZigBee, which has a mesh core of full function nodes, whilst edge nodes are reduced function, not being redundantly routed. This is a cost-benefit trade-off; edge nodes may be designed to be less complex since they do not have to participate in routing. The drawbacks include the facts that redundancy and hence reliability is reduced, and the ability to extend the network beyond the edge, in an ad hoc manner, is lost.

Note that a fully meshed network is not necessarily a fully connected network, in the strict sense. The fully connected network has a single hop to all other nodes. A full mesh has single or multiple hops to all other nodes and redundant route potential at all nodes.

Where its detail is already familiar to the reader, it may helpful to recall that the Internet is a mesh. However, in the absence of load balancing or link aggregation approaches, it often works simply as a per-packet or per-flow, dynamically reconfigured, multi-hop network.

Returning to the access mesh, this always tends towards multi-hopping since the presumption is often that the access node is the gateway for most traffic, creating a natural tree and branch structure.

1.4 Key mesh issues and the structure of this book

This chapter now moves on to list the key issues for mesh networking. The issues are briefly introduced here and each will be subsequently followed up in the succeeding chapters of this book. It is convenient to introduce the issues via a series of questions one may want to ask about meshes. In fact, the questions are real, having previously been gathered by interviewing networking industry professionals in 2005.

The first question is simply 'What are meshes good for?' The answers are best given by application examples, which include

- cellular or WLAN hot-spot multi-hopping,
- community networking,
- home and office indoor networking,
- micro base station backhaul,
- vehicular ad hoc networks (VANETs), and
- wireless sensor networks (WSNs).

These six application examples are described at the beginning of Chapter 2, with deeper consideration of the fundamental technology following in Chapter 3.

We found that the next questions typically asked about meshes included one or all of the following.

- Do meshes improve spectral efficiency ?
- Do meshes self-generate capacity ?
- Do directional antennas help a mesh?
- Can meshes improve the overall utilisation of spectrum?

These four questions come up so commonly, and so commonly together, that we shall spend quite some time looking at them. In fact these questions are so core to understanding meshes that we shall elevate them all to the level of hypotheses in this book. We can then test them formally in Chapter 4, to determine whether they are true or false, by examining all the available evidence.

Next, having satisfied ourselves as to the validity or otherwise of the four hypotheses, we shall begin to look more closely at meshes in practical deployments. It transpires that one question a seasoned network operator would be right to ask is 'Are meshes more susceptible to interference?' Here the answer really is uncertain, since it depends on the detail of implementation. This is because if the mesh does encounter interference, it may re-route to a better path. But unfortunately this re-routing, especially if it occurs frequently, may upset the transport layer of the communications protocol, potentially leading to overall worse performance. We look at mesh susceptibility in Chapter 5.

The network operator may continue to ask 'Are there any unusual aspects of mesh networking?' We shall see the answer is indeed 'Yes'. For example, in a mobile mesh users may realise their battery is constantly being drained even if they are not actively using it. Looking back to the house coverage example earlier in this chapter, this would apply to the third person's node, which is relaying traffic to/from the end of the garden. This user might decide to turn off their node, with undesirable consequences for the whole mesh connectivity. A related point, but this time from the point of view of an operator, is the realisation that user connectivity is no longer under the operator's sole control. We shall see that there are some steps which can be taken to mitigate this effect. This book tackles the subject by first looking at mesh routing mechanisms in Section 5.3 as a foundation for quality of service issues in Chapter 6.

One question the reader may well have by now is 'Does this book offer any hints and tips for a good mesh deployment?' After all the main technical arguments have been developed, we include Chapter 7 which reviews and collects the common pitfalls to be avoided when deploying mesh networks. Inside the back cover of the book, we list a concise summary of hints and tips.

The six application examples given above are also revisited near the end of the book in Chapters 8 and 9, so that all the theory developed in the