1

Introduction

One traditional view of how wireless networks evolve is of a continuous, inevitable progression to higher link speeds, combined with greater mobility over wider areas. This standpoint certainly captures the development from first and second generation cellular systems focused on voice support, and the early short-range wireless data networks, through to today's 3G cellular and mobile broadband systems; there is every confidence that the trend will continue some way into the future. Such a picture neatly summarizes a massive body of research and development of radio technologies, from antenna design to link coding to radio resource optimization. Pictures such as Figure 1.1 are well known from discussions of future wireless systems.

However, this book takes quite a different perspective. Instead of starting from the physical layer problems of wireless systems, the focus is on the networking issues that arise as the communications world moves towards offering ever more sophisticated services in more complex commercial environments. Furthermore, although these questions arise most prominently for the mobile and wireless domain – partly because the very nature of wireless communication encourages diversity in the business relationships and partly because the technical challenges require diversity in the physical layer solutions – the same issues will arise in any networking context. The resulting trend towards increasing technological and administrative heterogeneity is the one which has to be addressed primarily at the network level, and the pressures that it causes may lead communications systems to look radically different in the future from how they look today. This book presents a snapshot of current research into a set of new networking concepts that will enable such a vision.

1.1 The Current Communications Environment

There are already successful standards for mobile and wireless networking which fully address today's markets and existing air interfaces. The standardization bodies responsible for cellular systems, primarily the 3rd Generation Partnership Project (3GPP) [1], and for data networks, primarily the IEEE 802 LAN/MAN Standards Committee [2], both maintain their current standard systems and have a continuous programme of enhancements and

Ambient Networks: Co-operative Mobile Networking for the Wireless World Norbert Niebert (Ericsson GmbH), Andreas Schieder (Ericsson GmbH), Jens Zander and Robert Hancock

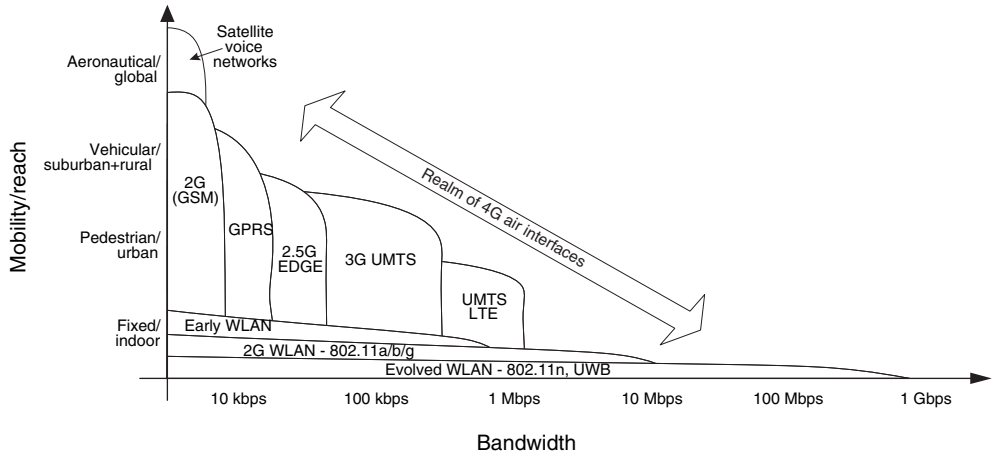

Figure 1.1 Trends in the wireless physical layer

system evolution. For example, 3GPP networks are already rolling out high-speed uplink and downlink packet access (HSDPA/HSUPA) as an extension of the current third-generation air interface. At the same time, they are working on a long-term evolution for the radio access network, and the evolution of the system architecture in general, activities referred to as LTE/SAE respectively [3].

This book is not primarily about these new developments *per se*; they can be seen as relatively low-risk incremental changes to current architectures and deployed networks, as is quite appropriate work for standards bodies to undertake. Rather, they are examples of the increasing complexity that will eventually require a new way of thinking about the way the mobile communications networks are put together. This growth in complexity is actually the result of two more fundamental, underlying trends.

Foremost among these is a change in the business environment. The starting point has been a vertically integrated model, where a complete end-to-end service, including access provision and infrastructure management, is provided by a small number of operators, supplemented by international roaming. The trend is towards much more complex models. The first aspect of this is a lengthening and fragmentation of the traditional value chain, allowing entities to focus on and specialize in particular activities such as service creation and marketing, or infrastructure operation. This already happens in the cellular market, where a set of new interoperator interfaces has had to evolve to support it. The same trend is visible in the integration of new access technologies such as WLAN, where service and access provision are almost invariably split in the 3GPP interworking case, bringing the additional complexity of the need to offer the same services over radically different bearer types. The rise of the hotspot market also presents new scaling problems, as the number of individual operators is larger by some orders of magnitude compared to the cellular world. Finally, to maintain growth, there is a need for the mobile world to extend to embrace new communications markets – not just the enterprise, but also home and personal networks. Along with the issues of scaling and heterogeneity already mentioned, any such development will create further difficulties for internetworking: the relatively open and unmanaged nature of the environment and the wide variation in

business models will mean that traditional forms of interoperator agreement will no longer be sufficient.

The second major motivation is the rate of change in the technological environment, a rate which shows no sign of decreasing. Along with the introduction of new air interfaces (mentioned above) there is also evolution, driven by basic economic and engineering requirements, in the configurations in which access infrastructure needs to be deployed. Examples here are vehicular networks, which can insulate user terminals from the special problems of high physical mobility, and meshed wireless networks, which reduce the cost of achieving area coverage with very short range air interfaces. Along with encouraging the business evolution referred to above, these developments present challenges for existing systems as they cannot be reconciled with the assumptions about air interface behaviour or functionality distribution that are implicit in the network architectures. At the other end of the protocol stack, there is similar if not more rapid change in the range of services that networks are expected to accommodate. The changes encompass both the type of service (from voice to data and multimedia) and the users (extending to peer–peer operation). These developments make additional demands on the flexibility of the network in the efficient mapping of services to very heterogeneous physical resources, and the routing and control of traffic within the networks. One common feature is the trend towards the Internet Protocol (IP) as a universal network layer, which is visible in both its use as the basis of these advanced services in 3GPP networks and its adoption even for wholesale replacement of fixed-line voice networks [4]. However, the core Internet standards do not offer the level of control that is necessary for the advanced scenarios that are being considered. Thus, when we consider either of these aspects of technological change, current system architectures are not able to adapt at the speed which the marketplace demands.

The combination of these trends – increasing heterogeneity at all levels and increasing demands for service complexity and control – can be seen as the networking counterpart to the physical layer trends that are more commonly used to mark out the evolution through the mobile network generations. The goals for the Ambient Networking concept can best be shown by an analogous picture, Figure 1.2.

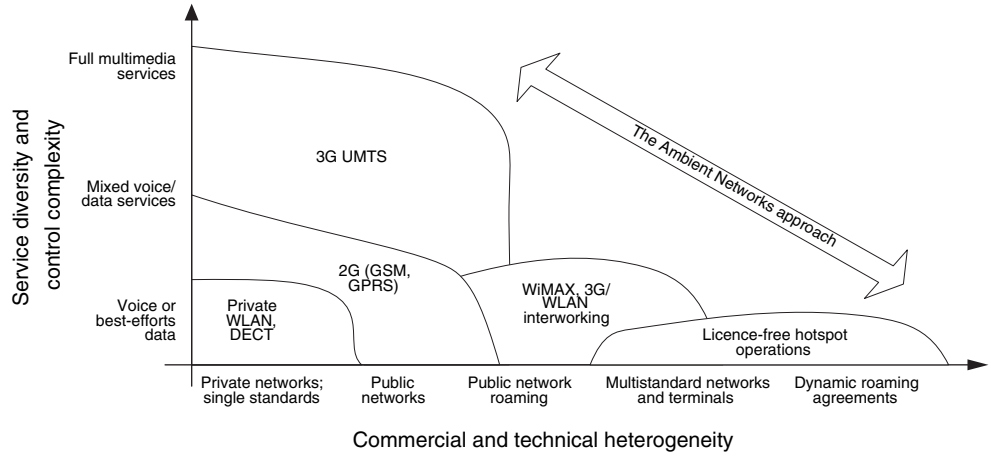

Figure 1.2 Networking aspects of system evolution

1.2 The Ambient Networking Concept

The Ambient Networking concept is a unification of a wide variety of new ideas from across the communications research community; however, it is not driven solely by scientific or technological goals. It also places a premium on developing solutions which are commercially exploitable, by taking into account the necessity to develop from current systems and building a consensus between the multiple different communities which make up the mobile world. In this sense, it certainly does not start from a clean sheet, but it does take the license to consider a more radical set of possibilities than would be encompassed by pure incremental evolution from current standards.

The core of the approach is the development of a set of control functions, which operate primarily at the network level. These functions can be implemented as an overlay on top of existing or new connectivity technologies, provided those technologies expose certain basic data transfer capabilities. This concentration on control functionality as an overlay is the key to addressing the twin problems of convergence between different technology types and migration from existing systems. Such an approach partly addresses the trend towards heterogeneity and service richness as described above, but introduces the risk of an explosion of complexity in network operation, especially as the commercial environment becomes more complex. The aim is to handle this issue by considering the set of control functions for a given network as a coordinated set. This goal is that this integrated Ambient Control Space will present a simpler interface to the outside world for external interworking. More ambitiously, given such a universal internetworking interface in the control plane, it should be possible to develop techniques whereby the control spaces of different networks can be recursively combined to support more complex scenarios in the same basic way. In the Ambient Networking context, this concept is referred to as *network composition*; it can be seen as the generalization to the control plane of the best efforts internetworking in the user plane (i.e. for forwarding and routing) that has been the foundation of the success of the Internet.

1.3 The Ambient Networks Project

It has been widely recognized for some time in the research community (see for example the references gathered by the NewArch project [5] and the final report [6]) that there is a need for new thinking to handle heterogeneity and control problems in the network layer and to enable the demands of different business actors to be arbitrated within a common technological framework. This book does consider trends in the research landscape in general, especially in particular technical areas such as resource management, context, overlay networks and so on. However, the core of the work, and the overall architectural concept, is based on the work done in the first phase of the collaborative project 'Ambient Networks' [7].

Ambient Networks is a joint industrial and academic research project. The industrial members come mainly from the network operator and equipment manufacturer communities, although other industries are also represented, following the view that networking functionality will eventually spread to all technology sectors. The academic members come from a wide variety of both universities and research institutes. The activity is sponsored under the 6th Framework Programme of the European Union as an 'Integrated Project' – in other words, a project which takes a set of related research activities and both develops them individually and integrates them into a unified conceptual vision. Indeed, this book has the same two-part

structure, describing an overall architecture and a set of specific technology solutions. As befits its origins, the consortium members are mainly based in Europe; however, the nature of the subject matter and the industry is that the perspective is entirely global, and the consortium also includes members from North America and the Asia/Pacific region. The first phase of the project involved over 150 different researches over two years of work, carrying out mainly conceptual investigations with initial simulations and demonstrations.

The Ambient Networks project should not be viewed in isolation. Historically, projects under previous EU Framework Programmes have been instrumental in setting the stage for major commercial developments, especially in the communications world. For example, in the 4th Framework, the FRAMES and RAINBOW projects (see e.g. [8,9]) provided the foundation for what rapidly became the air interface and network architecture for UMTS, and there are also links to current complementary work. In particular, although Ambient Networks focuses on network layer issues, it is recognized that there will be significant interactions with future air interface evolution and future terminal and service concepts. These are the subjects of sister projects within the 6th Framework Programme, under the general umbrella of the Wireless World Initiative (WWI). Further information on the WWI projects, and its associated open research community, the Wireless World Research Forum, can be found in [10,11].

1.4 How to Read This Book

The structure of this book follows the structure of the thinking behind the Ambient Networking approach itself: a conceptual framework and unifying architecture, within which a set of specific research topics is developed in more detail.

The first part of the book treats Ambient Networking at the overall level. Chapter 2 provides a technical perspective on the origins of the concept, in current thinking about convergence between different types of networks (fixed, mobile, wireless), treating the case of the introduction of the Internet Multimedia Subsystem in 3G networks as an example. From this starting point, the need for a new architectural approach and the requirements on that new architecture are derived. This is complemented in Chapter 3 by a discussion of the business perspective and economic drivers, including a description of a generic business model within which the Ambient Networks results can be analysed, with a particular focus on the issues that influence migration and deployment of new networking technologies. The architecture itself is presented in Chapter 4. The Ambient Networks architecture is in many respects deliberately minimalist, and the chapter begins with five basic principles from which most of the more specific architectural decisions have been derived. It then describes the two basic features of the architecture around which the details are arranged: the Ambient Control Space which provides an environment within which the various control functions are organized, and the Ambient Layer Model which captures how interactions with connectivity infrastructure and services and applications are codified. Most readers will find all of these three chapters relevant, albeit with a different level of importance depending on whether their focus is business or technical; the material of Chapter 4 is a prerequisite for following the remainder of the book.

The first part of the book continues with three chapters on specific technical aspects of the overall Ambient Networks concept. Chapter 5 describes the approach taken to security, starting with a survey of the problem space and assumptions about feasible security mechanisms, leading to a definition of the fundamental building blocks of the Ambient Networks security architecture – in particular, secure identification and authorization, and their application to

some specific security problems. Chapter 6 provides a detailed discussion of the network composition concept, first from a procedural perspective (how the composition process might actually take place) and then considering which types of composed networks might be produced as a result. Network composition is one of the key Ambient Networking concepts and implies a new set of requirements for creation and management of control relationships between networks; these control relationships will require support from a new family of signalling protocols. Chapter 7 presents the ambient signalling solution in the context of current IETF signalling protocols and concludes with a detailed example of the application of the signalling to the specific problem of internetwork service level agreement negotiation.

The second part of the book consists of five chapters which cover specific technical research which has been carried out within the Ambient Networks framework. These chapters can be read largely independently of each other and in any order, although in the book they are presented roughly in the sequence of the protocol stack. Chapter 8 presents work on multi-radio access, specifically the problems of integrating multiple different radio access technologies into a single system. There are two key concepts: an architecture for coordinating the resource management functions, and protocol components to unify a set of diverse link layers. A specific aspect is an analysis of the commercial benefits from multi-radio integration. Chapter 9 continues with a detailed study of the mobility management functions required in the network layer, again with a particular emphasis on methods for combining different mobility mechanisms in different network types. Chapter 10 considers the use of overlay network techniques to provide value-added functionality with the network infrastructure, to meet resource optimization requirements for media delivery, which are particularly critical in the wireless domain. The work includes detailed consideration of the scalability issues in management of large-scale overlay networks. Finally, Chapters 11 and 12 consider architecture for the integration of context information into network operation, including the definition of a common framework for context awareness across all functions in the Ambient Control Space, and the application of new ideas in network management to the Ambient Network environment.

1.5 Outlook

As we write these words, the material which this book describes is already being developed from its original conceptual form. The project itself has entered its second phase; here, the major emphasis is on formalizing the system architecture and its interface definitions and also on building a set of simulators and demonstration systems that can be used to show the concepts in action and quantify their benefits. At the same time, the first steps have been made in taking the work to the major standardization bodies, both for specific protocols and at the overall conceptual level. The Ambient Networking concept is itself evolving to meet the real challenges of implementation and deployment.