Preface

The field of computer networking has evolved significantly over the past four decades since the development of ARPANET, the first large-scale computer network. The Internet has become a part and parcel of everyday life virtually worldwide, and its influence on various fields is well recognized. The TCP/IP protocol suite and packet switching constitute the core dominating Internet technologies today. However, this paradigm is facing challenges as we move to next-generation networking applications including multimedia transmissions (IPTV systems), social networking, peer-to-peer networking and so on. The serious limitations of the current Internet include its inability to provide Quality of Service, reliable communication over periodically disconnected networks, and high bandwidth for high-speed mobile devices.

Hence, there is an urgent question as to whether the Internet's entire architecture should be redesigned, from the bottom up, based on what we have learned about computer networking in the past four decades. This is often referred to as the "Clean Slate" approach to Internet design. In 2005, the US National Science Foundation (www.nsf.gov) started a research program called Future Internet Network Design (FIND) to focus the research community's attention on such activities. Similar funding activities are taking place in Europe (FIRE: Future Internet Research and Experimentation), Asia, and other regions across the globe. This book is an attempt to capture some of the pioneering efforts in designing the Next-Generation Internet. The book is intended to serve as a starting point for researchers, engineers, students, and practitioners who wish to understand and contribute to the innovative architectures and protocols for the next-generation Internet.

Book organization

The book is divided into four parts that examine several aspects of next generation networks in depth.

Part I, titled "Enabling technologies," consists of five chapters that describe the technological innovations which are enabling the design and development of next-generation networks.

Chapter 1, "Optical switching fabrics for terabit packet switches," describes photonic technologies to realize subsystems inside high-speed packet switches and routers. The proposed architectures using optical interconnections remain fully compatible with current network infrastructures. For these architectures, the authors conduct scalability analysis and cost analysis of implementations based on currently available components.

Chapter 2, "Broadband access networks: current and future directions," describes Long-Reach Passive Optical Network (LR-PON) technology which brings the high capacity of optical fiber closer to the user. The authors propose and investigate the Wireless-Optical Broadband Access Network (WOBAN) which integrates the optical and wireless access technologies.

Chapter 3, "The optical control plane and a novel unified control plane architecture for IP/WDM networks," provides an overview of current protocols utilized for the control plane in optical networks. The authors also propose and investigate a new unified control plane architecture for IP-over-WDM networks that manages both routers and optical switches.

Chapter 4, "Cognitive routing protocols and architecture," describes the operation of wireless networks in which cognitive techniques are becoming increasingly common. The authors present cognitive routing protocols and their corresponding protocol architectures. In particular, the authors propose and investigate the mobility-aware routing protocol (MARP) for cognitive wireless networks.

Chapter 5, "Grid networking," describes Grid networks which are enabling the large-scale sharing of computing, storage, communication and other Grid resources across the world. Grid networks based on optical circuit switching (OCS) and optical burst switching (OBS) technologies are discussed. The authors describe approaches for resource scheduling in Grid networks.

Part II, titled "Network architectures," consists of five chapters that propose and investigate new architectural features for next-generation networks.

Chapter 6, "Host Identity Protocol (HIP)," describes a set of protocols that enhance the original Internet architecture by injecting a name space between the IP layer and the transport protocols. This name space consists of cryptographic identifiers that are used to identify application endpoints, thus decoupling names from locators (IP addresses).

Chapter 7, "Contract switching for managing inter-domain dynamics," introduces contract switching as a new paradigm for allowing economic considerations and flexibilities that are not possible with the current Internet architecture. Specifically, contract switching allows users to indicate their value choices at sufficient granularity, and providers to manage the risks involved in investments for implementing and deploying new QoS technologies.

Chapter 8, "PHAROS: an architecture for next-generation core optical networks," presents the Petabit/s Highly-Agile Robust Optical System (PHAROS), an architectural framework for future core optical networks. PHAROS, which is designed as part of the DARPA core optical networks (CORONET) program, envisions a highly dynamic network with support for both wavelength and IP services, very fast service setup and teardown, resiliency to multiple concurrent network failures, and efficient use of capacity reserved for protected services.

Chapter 9, "Customizable in-network services," proposes the deployment of custom processing functionality within the network as a means for enhancing the ability of the Internet architecture to adapt to novel protocols and communication paradigms. The chapter describes a network service architecture that provides suitable abstractions for specifying data path functions from an end-user perspective, and discusses technical challenges related to routing and service composition along the path.

Chapter 10, "Architectural support for continuing Internet evolution and innovation," argues that, while the current Internet architecture houses an effective design, it is not in itself effective in enabling evolution. To achieve the latter goal, it introduces the SILO architecture, a meta-design framework within which the system design can change and evolve. SILO generalizes the protocol layering concept by providing each flow with a customizable arrangement of fine-grain, reusable services, provides support for cross-layer interactions through explicit control interfaces, and decouples policy from mechanism to allow each to evolve independently.

Part III, titled "Protocols and practice," deals with different aspects of routing layer protocols and sensor network infrastructures.

Chapter 11, titled *Separating Routing Policy from Mechanism in the Network Layer*, describes a network layer design that uses a flat endpoint identifier space and also separates routing functionality from forwarding, addressing, and other network layer functions. This design is being studied as part of the Postmodern Internet Architecture (PoMo) project, a collaborative research project between the University of Kentucky, the University of Maryland, and the University of Kansas. The chapter also presents results from experimental evaluations, using a tunneling service that runs on top of the current Internet protocols.

Chapter 12, titled "Multipath BGP: motivations and solutions," discusses the motivation for using multi-path routing in the next generation Internet, in the context of the widely-used Border Gateway Protocol (BGP) routing protocol. The chapter then presents a set of proposed mechanisms that can interoperate with the existing BGP infrastructure. Solutions for both intra-domain and inter-domain multi-path routing are presented. These mechanisms are being implemented as part of the ongoing TRILOGY testbed, a research and development project funded by the European Commission.

Chapter 13, titled "Explicit congestion control: charging, fairness, and admission management," presents theoretical results on explicit congestion control mechanisms. These are a promising alternative to the currently used implicit congestion control mechanism of TCP. Examples of protocols using explicit congestion control are eXplicit Control Protocol (XCP) and the Rate Control Protocol (RCP). This chapter presents a proportional fair algorithm that can be used with RCP and an admission management algorithm that deals with the tradeoff between maximizing resource utilization and admission of burst arrivals.

Chapter 14, titled "KanseiGenie: a software infrastructure for resource management and programmability of wireless sensor network fabrics," presents a software framework that allows a community of users to develop applications based on a network of deployed wireless sensor nodes. The framework allows the sensor nodes to be shared by multiple applications, using slicing and virtualization of the nodes' resources. This project has been implemented as part of the NSF GENI initiative and promises to change the way in which sensor networks will operate in the future.

Finally, **Part IV**, titled "Theory and models", deals with theoretical foundations and models, as applicable to next generation Internet protocol design.

Chapter 15, "Theories for buffering and scheduling in Internet switches," presents interesting theoretical results on the use of small buffer sizes in the design of next generation routers/switches. It presents results on the interactions between a router's buffer size and TCP's congestion control mechanisms; and results on queueing theory based analysis on the fluctuation of traffic arrivals at a router. Based on these results, an active queue management mechanism is presented.

Chapter 16, "Stochastic network utility maximization and wireless scheduling," discusses network utility maximization (NUM) as a refinement of the layering as optimization decomposition principle that is applicable to dynamic network environments. The chapter provides a taxonomy of this research area, surveys the key results obtained over the last few years, and discusses open issues. It also highlights recent progress in the area of wireless scheduling, one of the most challenging modules in deriving protocol stacks for wireless networks.

Chapter 17, "Network coding in bi-directed and peer-to-peer networks," examines the application of network coding principles to bi-directed and peer-to-peer (P2P) networks. With the increasing use of P2P networks, it is essential to study how well network coding can be useful in such networks. The chapter discusses fundamental limitations of network coding for such networks and derives performance bounds. For P2P networks, the chapter presents practical network coding mechanisms for peer-assisted media streaming and peer-assisted content distribution.

Chapter 18, "Network economics: neutrality, competition and service differentiation," argues that the current Internet is not living up to its full potential because it delivers insufficient or inconsistent service quality for many applications of growing importance. The author explores how pricing can help expose hidden externalities and better align individual and system-wide objectives by structuring payments between content providers, ISPs, and users, to create the right incentives. The role of service differentiation in remedying problems that arise when users have heterogeneous requirements or utility functions is also examined.

Acknowledgements

We are grateful to the contributing authors for their efforts and diligence. This book would not have been possible without their help and cooperation – our sincere thanks to all for their contributions.

We also acknowledge the National Science Foundation which has provided partial support for our work on this project under grant CNS-0626741.

We thank Cambridge University Press for the opportunity to publish this book. We are especially grateful to Philip Meyler, Publishing Director for Engineering, Mathematical and Physical Sciences, Sabine Koch, Sarah Matthews, Sarah Finlay and the Cambridge TEXline staff for their guidance throughout this task. We also like to thank Ms. Yuyan Xue (UNL) who helped us in putting together the book materials. Thanks are also due to Google, Inc. for the *Google Docs* and *Google Sites* services, which facilitated cooperation and communication among the book editors and authors. We also acknowledge our respective institutions for providing us with the necessary computing infrastructure.