

To my parents
and Gabriele

Foreword

During the last decade, the need for realizing communication in networks has grown enormously. The use of telecommunication networks, local area computer networks, or wide area networks like the Internet is growing rapidly, and the current trend to integrate different types of data like text, audio, or video is further increasing the demand for high bandwidth, low latency communication networks. Similarly, tightly coupled parallel computer systems demand sophisticated communication devices.

The importance of realizing communication in networks has motivated intensive worldwide research activities, also in basic computer science research. Many proposals have been made for the topology of communication networks: meshes, hypercubes, hypercube-like networks such as the butterfly, shuffle exchange, de Bruijn and fat tree, and expanders. Many routing modes have been proposed: circuit switching, store-and-forward routing, wormhole routing, etc. Distinctions between oblivious and adaptive routing have been introduced, and hardware restrictions like a limited buffer size or edge-bandwidth have been taken into consideration. A large collection of routing protocols has been developed and analyzed for the networks and modes mentioned above.

This monograph provides a comprehensive description of the current research on routing. In particular, it presents essential new contributions on universal routing. By introducing the routing number of a network, i.e., the offline routing time for worst case permutations, Christian Scheideler offers a rigorous approach to measure the quality of routing protocols. He applies it to known protocols like Ranade's random rank protocol, its variants for bounded and unbounded buffers, and extensions to arbitrary networks. The main contributions are new universal protocols for store-and-forward and wormhole routing with small buffers and without buffers (deflection routing). He examines the benefits of large edge-bandwidth, and the capabilities and limitations of deterministic protocols and of bounded storage capabilities of the switches.

This monograph is an extension of Christian Scheideler's dissertation (Ph.D. thesis), submitted to Paderborn University in 1996. I am deeply impressed by his intuition about combinatorial and probabilistic phenomena, and by his capabilities to understand, apply, and extend deep combinatorial and probabilistic methods.

November 1997

Friedhelm Meyer auf der Heide
Paderborn University, Germany

Preface

Routing strategies are methods that use the available infrastructure between processing units to enable the exchange of information between them. The main components of routing strategies form methods for the selection of routes in a communication network and methods for scheduling the motion of the messages along their selected routes. Efficient routing strategies are of fundamental importance for two reasons.

We live in an era that is often referred to as the Information Age. Analyzing, manipulating, gathering, and distributing information has become an important economic and therefore political factor. Telecommunication networks, computer networks in companies and universities or the Internet are examples of networks that have been designed for a fast exchange of information and execution of requests (for instance, telephone calls, emails, money transfers between banks). The vast growth of users and services especially in metropolitan and wide-area networks makes it necessary to exploit the available performance of these networks as well as possible.

On the other hand, there are many scientific problems, such as modeling global weather patterns, analyzing the aerodynamic properties of a wing, and simulating the strange subatomic world of quantum theory, that require enormous volumes of computation, which cannot be handled by single processor systems. Hence multiprocessor machines are needed that can efficiently execute parallel programs that serve these purposes. Since the time for one computation step is usually much shorter than the time needed for exchanging information, much work has been invested in recent years to speed up the communication time as much as possible.

Many routing strategies have already been developed for specific classes of network topologies. Of special interest are so-called *universal* routing strategies, that is, strategies that can be applied to arbitrary network topologies. In addition to providing a unified approach to routing in standard networks, the advantage of universal routing strategies is that they are ideally suited for communication in irregular networks that are used in wide-area networks and that arise when standard networks are modified or develop faults. Furthermore, universal routing places no restrictions on the pattern of communication that is being implemented (such as requiring that it form a per-

mutation). Hence these algorithms are ideally suited for any communication problem that may occur during the execution of a parallel algorithm.

Since there have been significant advances in the development and analysis of universal routing strategies only during the last few years, almost none of the results have found their way into textbooks yet. The purpose of this monograph, which is based on my thesis written in 1996 at the Paderborn University, is therefore to present the history and state of the art of universal routing strategies.

The text is self-contained with respect to the historical background and notations used. However, it is recommended that the reader has some experience with using probabilistic arguments in proofs.

Chapter 1 gives an introduction to the area of communication strategies and presents research areas to which routing has a close relationship. In Chapter 2 we motivate our routing models by describing how routing is done in practice. Chapter 3 contains all terminology needed for the following chapters. After giving some basic definitions in probability theory and graph theory, we define a hardware model for parallel systems with point-to-point communication, the routing problem, routing strategies, models for passing messages, and models for storing routing information. The rest of this book gives a survey on universal routing strategies for the two most important message passing models studied in theory: the store-and-forward routing model and the wormhole routing model.

Chapters 4 to 9 deal with store-and-forward routing strategies. Chapter 4 gives a survey of the history of store-and-forward routing protocols. It presents both results about protocols for specific networks and universal routing protocols. Furthermore, networks are presented for which optimal randomized and deterministic store-and-forward routing protocols are already known. Chapter 5 introduces an important parameter for measuring the routing performance of networks. This parameter is called the *routing number*. The nice property of this parameter is that it describes the routing performance of networks more accurately than other parameters that have been used so far, such as the expansion or bisection width of a network. Furthermore, it can be easily used to demonstrate how close the performance of routing protocols can get to the optimal routing performance of networks. In Chapter 6, we prove the existence of efficient offline protocols for routing packets along a fixed path collection and show how these protocols can be applied to network simulations. The proofs for the first two offline protocols concentrate on minimizing the routing time, whereas the proof for the third offline protocol concentrates on minimizing the available space for buffering packets during the routing. Chapter 7 gives an overview of the best universal oblivious protocols for store-and-forward routing known so far. It is shown how each of these protocols can be applied to routing in specific networks, and what the limitations of these protocols are. In Chapter 8, a historical survey of adaptive routing protocols is given. Furthermore, some techniques

for developing universal adaptive protocols are presented. One of these techniques is called “routing via simulation”. It will be demonstrated how this strategy can be used to construct efficient deterministic routing protocols for arbitrary networks. Chapter 9 gives an overview of compact routing strategies, that is, strategies for networks in which the space for storing packets and routing information is limited. After summarizing results about how space restrictions influence the efficiency of selecting routes, we present trade-offs between space requirements and the time necessary to route messages in arbitrary networks. The “routing via simulation” technique presented in the previous section is generalized to a strategy that yields efficient universal compact routing protocols.

Chapters 10 to 12 deal with wormhole routing strategies. Chapter 10 gives a survey of the history of wormhole routing protocols. It presents both, results about protocols for specific networks, and results about universal routing protocols. Furthermore, upper and lower bounds are shown for wormhole routing in arbitrary networks. In Chapter 11 we describe two universal oblivious wormhole routing protocols and show how they can be applied to arbitrary networks. We further show how to improve the performance of one of these protocols for specific classes of networks such as butterflies and meshes. Chapter 12 deals with all-optical routing. We present nearly tight bounds for the runtime of a very simple protocol for sending messages along an all-optical path collection. We further show how this protocol can be applied to specific networks.

The monograph ends with a summary of its results and important open problems in the field of routing theory. Also, an outlook on future directions will be given.

Large parts of this book have been previously published as joint work. The results of Chapter 5 partially extend results in [MS96b]. The idea behind the third offline protocol in Chapter 6 is based on [MS95b]. Chapter 7 contains results that partially extend [MS95a, CMSV96]. Large parts of Chapters 8 and 9 are extensions of [MS95a, MS96a]. Chapter 11 is mostly based on [CMSV96], and Chapter 12 consists of results in [FS97].

Acknowledgements

First of all, I would like to thank Prof. Dr. Friedhelm Meyer auf der Heide for his great support and many helpful discussions. I am obliged to him for helping me writing the papers that led to this book and for establishing many important national and international contacts that influenced my work. Furthermore, I would like to thank Berthold Vöcking for the very productive collaboration and many valuable discussions. I also want to express my thanks to all members of the research group Meyer auf der Heide for the good cooperation and the nice atmosphere.

Besides my colleagues, I am very grateful to my parents whose continuous support greatly helped me to finish my diploma and this thesis. Last, but not least, I thank my wife Gabriele for her great support and for enduring all the days in which she hardly ever saw me because of my work.

Paderborn, November 1997

Christian Scheideler

This work has been partly supported by

- DFG-Sonderforschungsbereich 376 “Massive Parallelität: Algorithmen, Entwurfsmethoden, Anwendungen”,
- DFG Leibniz Grant Me872/6-1, and by
- EU ESPRIT Long Term Research Project 20244 (ALCOM-IT)