"So, why is your thesis my homework problem?"

— Grumblings of a Tired Graduate Student[†]

Part II: A Note to the Reader

In Part I of the thesis, we considered the overall router architecture. The switch fabric in any router architecture receives data from all N ports of a router. The total memory access rate and bandwidth handled by the switch fabric are proportional to NR. We analyzed various switch fabrics and router architectures, and described several load balancing and caching algorithms in Chapters 2-6 to alleviate this memory performance bottleneck.

In the second part of the thesis, we will consider the design of router line cards. A line card may terminate one or more ports in a router; thus the rate at which packets arrive on a line card is the aggregate of the line rates of all the ports that are terminated by it. We will denote R to be this aggregated line rate. Indeed, there are routers where the data from all ports in the router is aggregated onto one line card, making the memory access rate and bandwidth requirements as stringent as the switch fabric designs we encountered in Part I.

Unlike switch fabrics, whose main job is to transfer packets, the tasks on a router line card are of many different types, as described in Figure 1.6. We will only consider those tasks on the line card for which memory is a significant bottleneck, as described in Chapter 1. These tasks include packet buffering (Chapter 7), scheduling (Chapter 8), measurement counters (Chapter 9), state maintenance (Chapter 10), and multicasting (Appendix J).

In the rest of this thesis, we will separately consider each of these memory-intensive tasks. Since the aggregated data rate on a line card is R, and since R is a variable, the

[†]EE384Y, Stanford University, May 2002.

designs that we derive can be scaled to any line rate, and for line cards that terminate one or more ports. For example, a line card design with $R = 100$ Gb/s can be used to built a system with one 100 Gb/s port, ten 10 Gb/s ports, or a hundred 1 Gb/s ports.

We have implemented [33] the approaches described in Part II extensively in industry. In the rest of the chapters, we will give real-life examples of their use.[10] Indeed, as noted above, the packet processing ASICs and memories on line cards built in industry can support any combination of ports and line rates demanded by the customer.

In the course of implementation, some of the techniques and algorithms that we present in this thesis have been modified. Changes were also made where necessary to adapt them for specific market requirements and applications, in addition to the ones described here. Chapter 11 summarizes these observations and describes some remaining open questions.

[10] Our examples are from Cisco Systems [133].