DESIGNING A PREDICTABLE BACKBONE NETWORK USING VALIANT LOAD-BALANCING

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Rui Zhang-Shen
June 2007

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Nick McKeown) Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Balaji Prabhakar)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Ramesh Johari)

Approved for the University Committee on Graduate Studies.

To my parents,
who taught me the most important lessons.

Abstract

The backbones form the core of the Internet and carry large volumes of data across long distances. Designing a backbone network with performance guarantees is important but difficult due to the uncertainty in the traffic matrix, the need to accommodate component failures, and the desire to avoid congestion. Backbone networks today are highly over-provisioned but cannot guarantee to support all traffic matrices.

In this thesis we propose using Valiant Load-Balancing (VLB, named after L.G. Valiant) to design backbone networks that can support all traffic matrices efficiently, even under a number of failures.

Chapter 2 shows that VLB is efficient for any network. A simple scheme, the Gravity Full Mesh, requires a total capacity no more than twice the theoretical lower bound, and another scheme, Minimum Network Fanout, requires a total capacity at most 20% more than the theoretical lower bound.

The path diversity in VLB enables the network to tolerate failures efficiently. Chapter 3 shows that in order to tolerate k *arbitrary* failures, the fraction of extra capacity required is only approximately $\frac{k}{N}$. In addition, VLB provides fast rerouting after failure.

In Chapter 4 we first argue that most applications will not be affected by the extra propagation delay caused by VLB. Then we propose adaptive VLB, which adjusts the amount of load-balancing according to the traffic condition in order to minimize packet delay.

In Chapter 5 we propose using VLB to route traffic between two networks. VLB can efficiently utilize the peering links so that there is no congestion unless the total

peering traffic rate exceeds the total peering capacity. It can further ensure that peering packets receive the same quality of service as non-peering traffic. We derive the optimal load-balancing parameters for routing local traffic and peering traffic.

We finally apply VLB to designing circuit-switched networks with performance guarantees. Current circuit-switched networks use heuristics to minimize average blocking probability. The heuristics require parameters that depend on the traffic metrics, are hard to analyze, and cannot give performance guarantees. Valiant Load-Balancing, on the other hand, can give theoretical bounds on blocking probabilities for all flows under worst case traffic.

Acknowledgments

It has been a great privilege to work with my adviser Nick McKeown. Nick is a teacher, a mentor, and a role model. On research he is visionary; on everything else he is extremely understanding. Nick is always there to support, to guide, and to help, with his calming influence and amazing wisdom. I have learned so much from you, and thank you, Nick.

My undergraduate adviser Doug Osheroff is another amazing person I met at Stanford. Thank you, Doug, for inviting me to join Stanford and to join your lab, for guiding me, and for supporting me.

I am lucky to have worked with such great professors as Balaji Prabhakar and Ramesh Johari. Thank you for what you taught me and for being the readers of this thesis. I am also lucky to have worked with some top researchers and engineers outside of Stanford, especially John Chuang, Nasser El-Aawar, Sandy Fraser, Albert Greenberg, Frank Kelly, Murali Kodialam, TV Lakshman, Darren Loher, Jennifer Rexford, and Jennifer Yates. I thank them for what I learned from them.

My special thanks go to Nandita Dukkipati, my office mate and best friend. My graduate school years would not have been the same without you and your support. I will always remember our passionate discussions, work related or not. I also thank the other members of my research group, especially Clay, Da, David, Gireesh, Glenn, Guido, Isaac, Jad, Jianying, John, Justin, Martin, Masa, Neda, Pablo, Pankaj, Sundar, and Yashar, and members of Balaji's group, especially Abtin, Chandra, Damon, Devavrat, Elif, Kostas, Lu Yi, Mei, Paolo, and Rong. I enjoy all the meetings and discussions we had.

I am grateful to the Lillie Family Stanford Graduate Fellowship, the National

Science Foundation, and the Stanford Childbirth Accommodation Fund for their financial support.

I thank Marcia Keating and Judy Polenta for making my life at Stanford a much easier one than it could be, and Dr. Evelin Sullivan for editing part of this thesis.

My gratitude also goes to my earlier teachers, especially Yousheng Shu, Xun Yang, Zengcai Liu, and Xiaobin Li, who not only taught me but also encouraged me to pursue my dreams.

Last but not least, my deepest gratitude goes to my family: my grandparents, my parents, my husband, and my son. You all taught me a lot and your love and support define who I am. I am proud of you all as you are of me.

Contents

List of Tables

List of Figures